NVIDIA Onyx User Manual v3.10.4504 LTS

Table of Contents

1 Overview 19
1.1 Intended Audience 19
1.2 Related Documentation 19
1.3 Terminology 19
1.4 System Features 21
1.5 Ethernet Features 22
2 Getting Started 25
2.1 Configuring the Switch for the First Time 25
2.1.1 Configuring the Switch with ZTP 30
2.1.2 Rerunning the Wizard 30
2.2 Starting the Command Line (CLI) 30
2.3 Starting the Web User Interface (WebUI) 31
2.4 Zero-touch Provisioning 33
2.4.1 Running DHCP-ZTP 33
2.4.2 ZTP and OS Upgrade 35
2.4.3 DHCPv4 Configuration Example 35
2.4.4 DHCPv6 Configuration Example 35
2.4.5 ZTP Commands 36
3 User Interfaces 38
3.1 LED Indicators 38
3.2 Command Line Interface (CLI) 38
3.2.1 CLI Modes 38
3.2.2 Syntax Conventions 39
3.2.3 Getting Help 40
3.2.4 Prompt and Response Conventions 40
3.2.5 Using the “no” Command Form 41
3.2.6 Parameter Key 42
3.2.7 CLI Pipeline Operator Commands 43
3.3 Secure Shell (SSH) 46
3.3.1 Adding a Host and Providing an SSH Key 46
3.3.2 Retrieving Return Codes When Executing Remote Commands 47
3.4 Web Interface Overview 47

3.4.1 Password Hardening 47
3.4.2 Changing Default Password 48
3.4.3 About Web UI 49
3.4.4 Setup Menu 50
3.4.5 System Menu 51
3.4.6 Security Menu 51
3.4.7 Ports Menu 52
3.4.8 Status Menu 52
3.4.9 ETH Mgmt Menu 53
3.4.10 IP Route Menu 54
3.5 UI Commands 54
3.5.1 CLI Session 54
3.5.2 Web Interface 73
4 System Management 81
4.1 Management Interfaces 81
4.1.1 Configuring Management Interfaces with Static IP Addresses 81
4.1.2 Configuring IPv6 Address on the Management Interface 81
4.1.3 Dynamic Host Configuration Protocol (DHCP) 82
4.1.4 Default Gateway 82
4.1.5 In-Band Management 82
4.1.6 Configuring Hostname via DHCP (DHCP Client Option 12) 83
4.1.7 Management VRF 84
4.1.8 Management Interface Commands 84
4.1.9 Control Plane Policing (CoPP) 109
4.2 Chassis Management 121
4.2.1 System Health Monitor 122
4.2.2 Power Management 123
4.2.3 Monitoring Environmental Conditions 124
4.2.4 USB Access 125
4.2.5 Unit Identification LED 125
4.2.6 System Reboot 126
4.2.7 Viewing Active Events 126
4.2.8 Chassis Management Commands 127
4.3 Management Source IP Address 141

4.3.1 ntp source-interface 141
4.3.2 Commands 141
4.4 Upgrade/Downgrade Process 154
4.4.1 Important Pre-OS Upgrade Notes 155
4.4.2 Upgrading Operating System Software 155
4.4.3 Upgrading HA Groups 157
4.4.4 Upgrading MLAG-STP Setup 158
4.4.5 Deleting Unused Images 159
4.4.6 Downgrading OS Software 159
4.4.7 Upgrading System Firmware 162
4.4.8 Image Maintenance Using ONIE 163
4.4.9 Software Management Commands 164
4.5 Configuration Management 171
4.5.1 Saving a Configuration File 171
4.5.2 Loading a Configuration File 171
4.5.3 Restoring Factory Default Configuration 171
4.5.4 Managing Configuration Files 172
4.5.5 Automated Periodic Configuration File Backup 173
4.5.6 Configuration Management Commands 174
4.6 Resource Scale 193
4.6.1 Resource Scale Commands 193
5 System Synchronization 195
5.1 NTP and Clock 195
5.1.1 NTP Authenticate 195
5.1.2 NTP Authentication Key 195
5.1.3 Additional Reading and Use Cases 196
5.1.4 NTP Commands 196
5.2 Precision Time Protocol (PTP) 206
5.2.1 PTP Principles 206
5.2.2 Clock Types and Operation Modes 207
5.2.3 PTP Domains 208
5.2.4 Securing PTP Infrastructure 211
5.2.5 Additional Reading and Use Cases 213
5.2.6 PTP Commands 213

5.3 Replace CRC with Timestamp 248
5.3.1 Main Functionality 249
5.3.2 Setup Configuration 249
5.3.3 Replace CRC with Timestamp Commands 251
6 Network Management Interfaces 253
6.1 SNMP 253
6.1.1 Standard MIBs 253
6.1.2 Private MIBs 253
6.1.3 Proprietary Traps 254
6.1.4 Configuring SNMP 255
6.1.5 Resetting SNMPv3 Engine ID 255
6.1.6 Configuring an SNMPv3 User 256
6.1.7 Configuring SNMP Notifications (Traps or Informs) 256
6.1.8 SNMP SET Operations 257
6.1.9 Additional Readings and Use Cases 263
6.2 JSON API 263
6.2.1 Authentication 263
6.2.2 Sending the Request 266
6.2.3 JSON Request Format 266
6.2.4 JSON Response Format 268
6.2.5 Supported Commands 270
6.2.6 JSON Examples 270
6.2.7 JSON Request Using WebUI 273
6.2.8 Additional Reading and Use Cases 275
6.3 Network Management Interface Commands 275
6.3.1 SNMP 276
6.3.2 JSON API 288
7 Virtualization 290
7.1 Limiting the Container's Resources 290
7.1.1 Memory Resources Allocation Protocol 290
7.1.2 CPU Resource Allocation Protocol 291
7.2 Upgrade Ramifications 291
7.2.1 Changing Docker Storage Driver 291
7.3 Additional Reading and Use Cases 292

7.4 Docker Containers Commands 292
7.4.1 docker 292
7.4.2 docker login 293
7.4.3 docker logout 293
7.4.4 commit 294
7.4.5 copy-sdk 294
7.4.6 remove image 294
7.4.7 exec 295
7.4.8 label 295
7.4.9 load 295
7.4.10 pull 296
7.4.11 save 296
7.4.12 shutdown 297
7.4.13 start 297
7.4.14 image upload 298
7.4.15 file image upload 299
7.4.16 show docker 299
7.4.17 show docker containers 300
7.4.18 show docker images 302
7.4.19 show docker ps 302
7.4.20 show docker labels 303
7.4.21 show docker login 303
7.4.22 show docker stats 303
8 Telemetry, Monitoring, and Debuggability 305
8.1 WHAT JUST HAPPENED 305
8.1.1 Configure What Just Happened (WJH) Using CLI 305
8.1.2 Configure WJH Using NEO 318
8.1.3 WJH Streaming and Integration with Telegraf, InfluxDB and Grafana (TIG) Stack 318
8.2 Logging 319
8.2.1 Monitor 319
8.2.2 Remote Logging 319
8.2.3 Logging Protocol 320
8.2.4 Logging Commands 320

8.3 Debugging 336
8.3.1 Additional Reading and Use Cases 337
8.3.2 Debugging Commands 337
8.4 Link Diagnostic Per Port 345
8.4.1 Link Diagnostic Commands 346
8.5 Signal Degradation Monitoring 347
8.5.1 Effective-BER Monitoring 348
8.5.2 Configuring Signal Degradation Monitoring 348
8.5.3 Signal Degradation Monitoring Commands 348
8.6 Event Notifications 349
8.6.1 Supported Event Notifications and MIB Mapping 350
8.6.2 Terminal Notifications 352
8.6.3 Email Notifications 352
8.6.4 Command Event Notifications 353
8.7 Port Mirroring 364
8.7.1 Mirroring Sessions 365
8.7.2 Configuring Mirroring Sessions 368
8.7.3 Verifying Mirroring Sessions 370
8.7.4 Additional Reading and Use Cases 370
8.7.5 Port Mirroring Commands 371
8.8 sFlow 375
8.8.1 Flow Samples 376
8.8.2 Statistical Samples 376
8.8.3 sFlow Datagrams 377
8.8.4 Sampled Interfaces 377
8.8.5 Configuring sFlow 377
8.8.6 Verifying sFlow 378
8.8.7 Additional Reading and Use Cases 378
8.8.8 sFlow Commands 378
8.9 Buffer Histograms Monitoring 383
8.9.1 Additional Reading and Use Cases 384
8.9.2 Buffer Histograms and Thresholds Commands 384
8.10 Statistics and Alarms 394
8.10.1 Commands 394

8.11 Management Information Bases (MIBs) 409
8.11.1 Calculating of entPhysicalIndex in the Entity MIB 409
8.11.2 Examples 411
9 Automation Tools 412
9.1 Ansible 412
9.2 SALT 412
9.2.1 Installing SaltStack on CentOS 7 413
9.2.2 Configuring Salt 413
9.2.3 Configuring the Salt-minion File 414
9.2.4 Configuring the Proxy 414
9.2.5 Creating the pillar Directory 414
9.2.6 Running Onyx Salt Commands on the Server 415
9.3 Scheduled Jobs 415
9.3.1 Commands 416
10 User Management, Authentication, & Security 421
10.1 User Management & Security 421
10.1.1 User Accounts 421
10.1.2 Authentication, Authorization, and Accounting (AAA) 421
10.1.3 User Re-authentication 422
10.1.4 RADIUS 422
10.1.5 TACACS+ 422
10.1.6 LDAP 422
10.1.7 System Secure Mode 423
10.1.8 User Management and Security Commands 425
10.1.9 802.1x Protocol 453
10.2 Cryptographic (X.509, IPSec) and Encryption 461
10.2.1 System File Encryption 461
10.2.2 Cryptographic and Encryption Commands 462
11 Quality of Service (QoS) 471
11.1 QoS Classification 471
11.1.1 Trust Levels 471
11.1.2 Switch Priority to IEEE Priority Mapping 471
11.1.3 Default QoS Configuration 472
11.1.4 Control Protocols 472

11.2 QoS Rewrite 473
11.2.1 Switch-priority to PCP,DEI Re-marking Mapping 473
11.2.2 Switch-priority to DSCP Re-marking Mapping 473
11.2.3 DSCP to Switch-priority in Router 474
11.2.4 Default Configuration 474
11.3 Queuing and Scheduling (ETS) 474
11.3.1 Traffic Class 474
11.3.2 Traffic Shapers 474
11.3.3 Default Shaper Configuration 475
11.4 RED and ECN 475
11.5 Additional Reading and Use Cases 476
11.6 QoS Commands 477
11.7 QoS Commands 477
11.7.1 QoS Classification 478
11.7.2 QoS Rewrite 488
11.7.3 Queuing and Scheduling (ETS) 491
11.7.4 RED & ECN 493
11.8 Priority Flow Control (PFC) 495
11.8.1 Flow Control Threshold Configuration 497
11.8.2 PFC Watchdog 497
11.8.3 Additional Reading and Use Cases 498
11.8.4 PFC Commands 498
11.9 Shared Buffers 502
11.9.1 Traffic Pool Configuration 503
11.9.2 Lossless Traffic 503
11.9.3 Advanced Buffer Configuration 504
11.9.4 Additional Reading and Use Cases 509
11.9.5 Shared Buffer Commands 509
11.9.6 Shared Buffer Commands 509
11.10 Storm Control 526
11.10.1 Storm Control Commands 527
11.11 Head-of-Queue Lifetime Limit 528
11.11.1 HoQ Commands 528
11.12 Store-and-Forward 529

11.12.1 Additional Reading and Use Cases 529
11.12.2 Store-and-Forward Commands 530
12 Ethernet Switching 531
12.1 Ethernet Interfaces 531
12.1.1 Breakout Cables 531
12.1.2 56GbE Link Speed 534
12.1.3 Transceiver Information 534
12.1.4 High Power Transceivers 534
12.1.5 Forward Error Correction 535
12.1.6 Port Recirculation 535
12.1.7 Ethernet Interface Commands 536
12.2 Interface Isolation 559
12.2.1 Configuring Isolated Interfaces 559
12.2.2 Interface Isolation Commands 561
12.3 Link Aggregation Group (LAG) 563
12.3.1 Configuring Static LAG 563
12.3.2 Configuring Link Aggregation Control Protocol (LACP) 564
12.3.3 Additional Reading and Use Cases 564
12.3.4 LAG Commands 565
12.4 Link Layer Discovery Protocol (LLDP) 576
12.4.1 Configuring LLDP 576
12.4.2 DCBX 577
12.4.3 Additional Reading and Use Cases 577
12.4.4 LLDP Commands 577
12.5 VLANs 585
12.5.1 Configuring Access Mode and Assigning Port VLAN ID (PVID) 585
12.5.2 Configuring Hybrid Mode and Assigning Port VLAN ID (PVID) 586
12.5.3 Configuring Trunk Mode VLAN Membership 586
12.5.4 Configuring Hybrid Mode VLAN Membership 587
12.5.5 Additional Reading and Use Cases 587
12.5.6 VLAN Commands 587
12.6 Voice VLAN 592
12.6.1 Configuring Voice VLAN 593
12.6.2 Limitations 595

12.7 Spanning Tree Protocol 595
12.7.1 Port Priority and Cost 595
12.7.2 Port Type 596
12.7.3 BPDU Filter 596
12.7.4 BPDU Guard 597
12.7.5 Loop Guard 597
12.7.6 Root Guard 598
12.7.7 MSTP 598
12.7.8 RPVST 598
12.7.9 STP Commands 600
12.8 MAC Address Table 615
12.8.1 Configuring Unicast Static MAC Address 615
12.8.2 MAC Learning Considerations 616
12.8.3 MAC Address Table Commands 617
12.9 MLAG 621
12.9.1 MLAG Keepalive and Failover 623
12.9.2 Unicast and Multicast Sync 624
12.9.3 MLAG Port Sync 624
12.9.4 MLAG Virtual System-MAC 624
12.9.5 Upgrading MLAG Pair 624
12.9.6 Interoperability with MLAG 626
12.9.7 Configuring MLAG 627
12.9.8 Additional Reading and Use Cases 632
12.9.9 MLAG Commands 632
12.9.10 MLAG Commands 632
12.10 Link State Tracking 643
12.10.1 Configuring Link State Tracking 643
12.10.2 Link State Tracking Commands 644
12.11 QinQ 646
12.11.1 QinQ Operation Modes 646
12.11.2 Configuring QinQ 647
12.11.3 QinQ Commands 648
12.12 Access Control List (ACL) 648
12.12.1 Configuring ACL 648

12.12.2 ACL Actions 649
12.12.3 ACL Logging 650
12.12.4 ACL Capability Summary 650
12.12.5 Additional Readings and Use Cases 653
12.12.6 ACL Commands 653
12.12.7 ACL Commands 653
12.13 User Defined Keys 693
12.13.1 Configuring UDK 693
12.13.2 UDK Commands 694
12.14 OpenFlow 696
12.14.1 Flow Table 697
12.14.2 OpenFlow 1.3 Workflow 698
12.14.3 Configuring OpenFlow 703
12.14.4 Configuring Flows Using CLI Commands 704
12.14.5 Configuring Secure Connection to OpenFlow 706
12.14.6 OpenFlow Commands 708
13 VXLAN 724
13.1 Configuring VXLAN 724
13.2 VMware Network Virtualization and Security Platform (NSX) Configuration 726
13.2.1 Hardware Topology 726
13.2.2 Switch Configuration 727
13.2.3 Adding the Switch to NSX 729
13.2.4 Mapping a Logical Switch to a Physical Switch Port 730
13.3 Additional Reading and Use Cases 731
13.4 RoCE Over VXLAN 731
13.4.1 RoCEv2 Using PFC and ECN 731
13.4.2 RoCEv1 Using PFC 732
13.5 VXLAN Commands 733
13.6 VXLAN Commands 733
13.6.1 protocol nve 734
13.6.2 interface nve 735
13.6.3 nve bridge 735
13.6.4 nve controller bgp 735
13.6.5 nve fdb flood bridge address 736

13.6.6 nve fdb flood load-balance 736
13.6.7 nve fdb learning remote 736
13.6.8 nve mode only 737
13.6.9 nve neigh-suppression 737
13.6.10 nve vlan bridge 738
13.6.11 nve vlan neigh-suppression 738
13.6.12 nve vni vlan 738
13.6.13 interface nve auto-vlan-map 739
13.6.14 interface nve disable nve vni 740
13.6.15 vxlan mlag-tunnel-ip 740
13.6.16 vxlan source interface loopback 740
13.6.17 shutdown 741
13.6.18 clear mac-address-table nve 741
13.6.19 clear nve counters 741
13.6.20 show interfaces nve 742
13.6.21 show interfaces nve detail 742
13.6.22 show interfaces nve counters 743
13.6.23 show interfaces counters vlan 743
13.6.24 show interfaces nve flood 744
13.6.25 show interfaces nve mac-address-table 744
13.6.26 show interfaces nve mac-address-table local learned unicast 745
13.6.27 show interfaces nve mac-address-table remote configured multicast 745
13.6.28 show interfaces nve peers 746
13.6.29 ovs ovsdb server 746
13.6.30 ovs ovsdb manager remote 747
13.6.31 ovs ovsdb server listen 747
13.6.32 ovs logging level 748
13.6.33 show ovs 748
14 Ethernet VPN (EVPN) 749
14.1 Overview 749
14.2 Example of How To Configure EVPN 750
14.2.1 Layer 2 Configuration, MLAG, and VLANs 751
14.2.2 Layer 3 Configuration 751
14.2.3 BGP and EVPN Configuration 753

14.2.4 Spine Configuration 754
14.3 Traffic Behavior During Failures 755
14.4 EVPN Troubleshooting 756
14.4.1 show interface nve 1 756
14.4.2 show interface nve 1 detail 756
14.4.3 show ip bgp evpn summary 757
14.4.4 show ip bgp evpn 757
14.4.5 show ip bgp evpn vni 10060 757
14.4.6 show ip bgp evpn with multiple filters 758
14.4.7 show mac-address-table 758
14.4.8 show ip arp 759
14.5 EVPN Data Center Interconnect (DCI) 759
14.5.1 Layer 2 DCI Connection 759
14.5.2 Layer 3 Routes WAN 760
14.6 EVPN Centralized L3 Gateway 760
14.6.1 Configuration Example of EVPN Centralized Gateway 760
14.6.2 Configuration Example of MLAG EVPN Centralized Gateway 762
14.7 EVPN Logging Examples 763
14.7.1 EVPN MAC Mobility Logs 763
15 IP Routing 764
15.1 IP Routing Overview 764
15.1.1 IP Interfaces 764
15.1.2 Equal Cost Multi-Path Routing (ECMP) 767
15.1.3 ARP Neighbor Discovery Responder 773
15.1.4 Policy Based Routing (PBR) 774
15.1.5 General IP Routing Commands 775
15.1.6 IPv6 823
15.2 OSPF 841
15.2.1 Router ID 842
15.2.2 ECMP 842
15.2.3 Configuring OSPF 843
15.2.4 Additional Reading and Use Cases 845
15.2.5 OSPF Commands 845
15.3 BGP 864

15.3.1 State Machine 864
15.3.2 Default Address Family 864
15.3.3 Default Route Originate 865
15.3.4 Peer Groups and Update Groups 865
15.3.5 Configuring BGP 865
15.3.6 Verifying BGP 866
15.3.7 Ethernet Virtual Private Network 867
15.3.8 BGP Unnumbered 867
15.3.9 Configuring BGP Unnumbered 868
15.3.10 Additional Reading and Use Cases 869
15.3.11 BGP Commands 869
15.3.12 BGP Commands 870
15.3.13 BGP Monitoring Protocol 928
15.4 Bidirectional Forwarding Detection (BFD) Infrastructure 931
15.4.1 Session Establishment 931
15.4.2 Interaction with Protocols 932
15.4.3 BFD Commands 932
15.5 Policy Rules 936
15.5.1 Route Map 936
15.5.2 Route Map Commands 937
15.6 VRRP 950
15.6.1 Load Balancing 951
15.6.2 Configuring VRRP 951
15.6.3 Additional Reading and Use Cases 954
15.6.4 VRRP Commands 954
15.7 MAGP 961
15.7.1 Configuring MAGP 961
15.7.2 Useful Reading and Use Cases 963
15.7.3 MAGP Commands 963
15.8 DHCP Relay 967
15.8.1 DHCP-R Virtual Routing and Forwarding (VRF) Auto-Helper 967
15.8.2 Upstream and Downstream Interfaces 968
15.8.3 DHCP Relay Commands 968
16 RDMA Over Converged Ethernet (RoCE) 983

16.1 RoCE Overview 983
16.1.1 Definitions/Abbreviation 983
16.2 Configuring RoCE 984
16.3 RoCE Commands 985
16.4 Further Information 985
16.5 RoCE Commands 986
16.5.1 roce 986
16.5.2 show roce 986
16.5.3 show interfaces ethernet 1/1 counters roce 989
16.5.4 clear roce interface ethernet 1/1 989
17 Multicast (IGMP and PIM) 991
17.1 Basic PIM-SM 991
17.2 Source-Specific Multicast (SSM) 992
17.3 Bidirectional PIM 992
17.4 PIM Load-Sharing 993
17.4.1 Rendezvous Point Load-Sharing 993
17.4.2 Next Hop Load-Sharing 994
17.5 Bootstrap Router 994
17.6 Configuring Multicast 994
17.6.1 Configuring IGMP 995
17.6.2 Verifying IGMP 995
17.6.3 Configuring PIM 996
17.7 Additional Reading and Use Cases 997
17.8 IGMP and PIM Commands 997
17.9 IGMP and PIM Commands 997
17.9.1 PIM 999
17.9.2 PIM Bidir 1013
17.9.3 Multicast 1017
17.9.4 IGMP 1021
17.10 IGMP Snooping 1027
17.10.1 Configuring IGMP Snooping 1028
17.10.2 Defining a Multicast Router Port on a VLAN 1028
17.10.3 IGMP Snooping Querier 1030
17.10.4 IGMP Snooping Querier Guard 1030

17.10.5 IGMP Snooping Commands 1031
18 Appendixes 1047
18.1 Appendix: Ethernet Storage Fabric (ESF) 1047
18.1.1 ESF Configuration using Ansible 1047
18.1.2 ESF Configuration Using CLI 1048
18.1.3 ESF Maintenance, Monitoring and Troubleshooting 1053
18.1.4 ESF Setup Examples 1060
18.2 Appendix: Enhancing System Security According to NIST SP 800-131A 1062
18.2.1 Web Certificate 1063
18.2.2 SNMP 1064
18.2.3 HTTPS 1065
18.2.4 Code Signing 1066
18.2.5 SSH 1066
18.2.6 LDAP 1066
18.3 Appendix: Feature Support per IC and CPU Type 1067
18.4 Appendix: Splunk Integration with NVIDIA Products 1068
18.4.1 Getting Started with Splunk 1068
18.4.2 Switch Configuration 1069
18.4.3 Adding a Task 1069
18.4.4 Retrieving Data from TCP and UDP Ports 1070
18.4.5 SNMP Input to Poll Attribute Values and Catch Traps 1072
18.5 Appendix: Show Commands Not Supported By JSON API 1076
18.6 Appendix: What Just Happened (WJH) Events 1077
19 Document Revision History 1079

This is a long-term support (LTS) release. LTS is the practice of maintaining a software product for an extended period of time (up to three years) to help increase product stability. LTS releases include bug fixes and security patches.

Welcome to NVIDIA Onyx Documentation

NVIDIA Onyx™ operating system enables the management and configuration of NVIDIA's Ethernet switch system platforms.

NVIDIA Onyx provides a full suite of management options, including support for SNMPv1, 2, 3, and web user interface (Web UI). In addition, it incorporates a familiar industry-standard CLI which enables administrators to easily configure and manage the system.

These pages provide information about the scope, organization, and command line interface of NVIDIA Onyx as well as configuration examples.

Software Download

To download the latest software, log in to the following website: enterprise-support.nvidia.com/s/

For common questions about the Enterprise Account please see the following webpage: nvid.nvidia.com/NvidiaUtilities/#/needHelp

Technical Support

Customers who purchased NVIDIA products directly from NVIDIA are invited to contact us through the following methods:

- URL: www.nvidia.com → Support
- E-mail: enterprisesupport@nvidia.com

Customers who purchased NVIDIA M-1 Global Support Services, please see your contract for details regarding Technical Support.
Customers who purchased NVIDIA products through an NVIDIA-approved reseller should first seek assistance through their reseller.

Document Revision History

A list of the changes made to the User Manual are provided in User Manual Revision History.

1 Overview

1.1 Intended Audience

These pages are intended for network administrators who are responsible for configuring and managing NVIDIA's switch platforms.

1.2 Related Documentation

The following table lists the documents referenced in this User Manual.

Document Name	Description
System Hardware User Manual	This document contains hardware descriptions, LED assignments, and hardware specifications, among other things
Switch Product Release Notes	Please look up the relevant switch system/series Release Notes file
Virtual Modular Switch Reference Guide	This reference architecture provides general information concerning NVIDIA L2 and L3 Virtual Modular Switch (VMS) configuration and design
Community	Provides Ethernet Switch Solutions

1.3 Terminology

Term	Description
AAA	Authentication, Authorization, and Accounting: • Authentication—verifies user credentials (username and password) • Authorization—grants or refuses privileges to a user/client for accessing specific services • Accounting—tracks network resources consumption by users
ARP	Address Resolution Protocol. A protocol that translates IP addresses into MAC addresses for communication over a local area network (LAN).
CLI	Command Line Interface. A user interface in which you type commands at the prompt.
DCB	Data Center Bridging
DCBX	Should be Data Center Bridging eXchange—an extension of Link Layer Data Protocol to discover DCB compliant peers and exchange configuration information
DHCP	The Dynamic Host Configuration Protocol (DHCP) is an automatic configuration protocol used on IP networks.
DNS	Domain Name System. A hierarchical naming system for devices in a computer network.
ECN	Explicit Congestion Notification
ETS	Enhanced Transmission Selection provides a common management framework for assignment of bandwidth to traffic classes.
FTP/TFTP/sFTP	File Transfer Protocol (FTP) is a standard network protocol used to transfer files from one host to another over a TCP-based network, such as the Internet.

Term	Description
Gateway	A network node that interfaces with both InfiniBand and Ethernet, using different network protocols.
HA	High Availability. A system design protocol that provides redundancy of system components, thus enables overcoming single or multiple failures in minimal downtime.
Host	A computer platform executing an Operating System which may control one or more network adapters
LACP	Link Aggregation Control Protocol (LACP) provides a method to control the bundling of several physical ports together to form a single logical channel. LACP allows a network device to negotiate an automatic bundling of links by sending LACP packets to the peer (directly connected device that also implements LACP).
LDAP	The Lightweight Directory Access Protocol is an industry standard application protocol for accessing and maintaining distributed directory information services over an IP network.
LLDP	Link Layer Discovery Protocol. A vendor neutral link layer protocol used by network devices to advertise their identify, capabilities and for neighbor discovery.
MAC	A Media Access Control address (MAC address) is a unique identifier assigned to network interfaces for communications on the physical network segment. MAC addresses are used for numerous network technologies and most IEEE 802 network technologies including Ethernet.
MTU	Maximum Transfer Unit. The maximum size of a packet payload (not including headers) that can be sent /received from a port.
Network Adapter	A hardware device that allows for communication between computers in a network.
NTP	Network Time Protocol. A protocol for synchronizing computer clocks in a network.
PFC/FC	Priority Based Flow Control applies pause functionality to traffic classes OR classes of service on the Ethernet link.
PTP IEEE-1588	Precision Time Protocol. A high-accuracy time transfer protocol for synchronizing computer clocks in a network.
RADIUS	Remote Authentication Dial In User Service. A networking protocol that enables AAA centralized management for computers to connect and use a network service.
RDMA	Remote Direct Memory Access. Accessing memory in a remote side without involvement of the remote CPU.
RoCE	RDMA over Converged Ethernet. A network protocol that leverages Remote Direct Memory Access (RDMA) capabilities to accelerate communications between applications hosted on clusters of servers and storage arrays.
RSTP	Rapid Spanning Tree Protocol. A spanning-tree protocol used to prevent loops in bridge configurations. RSTP is not aware of VLANs and blocks ports at the physical level.
SCP	Secure Copy or SCP is a means of securely transferring computer files between a local and a remote host or between two remote hosts. It is based on the Secure Shell (SSH) protocol.
SNMP	Simple Network Management Protocol. A network protocol for the management of a network and the monitoring of network devices and their functions.
SSH	Secure Shell. A protocol (program) for securely logging in to and running programs on remote machines across a network. The program authenticates access to the remote machine and encrypts the transferred information through the connection.
syslog	A standard for forwarding log messages in an IP network.
TACACS+	Terminal Access Controller Access-Control System Plus. A networking protocol that enables access to a network of devices via one or more centralized servers. TACACS+ provides separate AAA services.

1.4 System Features

Feature	Detail
Software management	• Dual software image • Software and firmware updates
File management	• FTP • TFTP • SCP
Logging	• Event history log • SysLog support
Management interface	• DHCP/Zeroconf • IPv6
Chassis management	• Monitoring environmental controls • Power management • Auto-temperature control • High availability
Network management interfaces	• SNMP v1,v2c,v3 • JSON
Security	• SSH • Telnet • RADIUS • TACACS+
Date and time	• NTP
Cables & transceivers	• Transceiver info

1.5 Ethernet Features

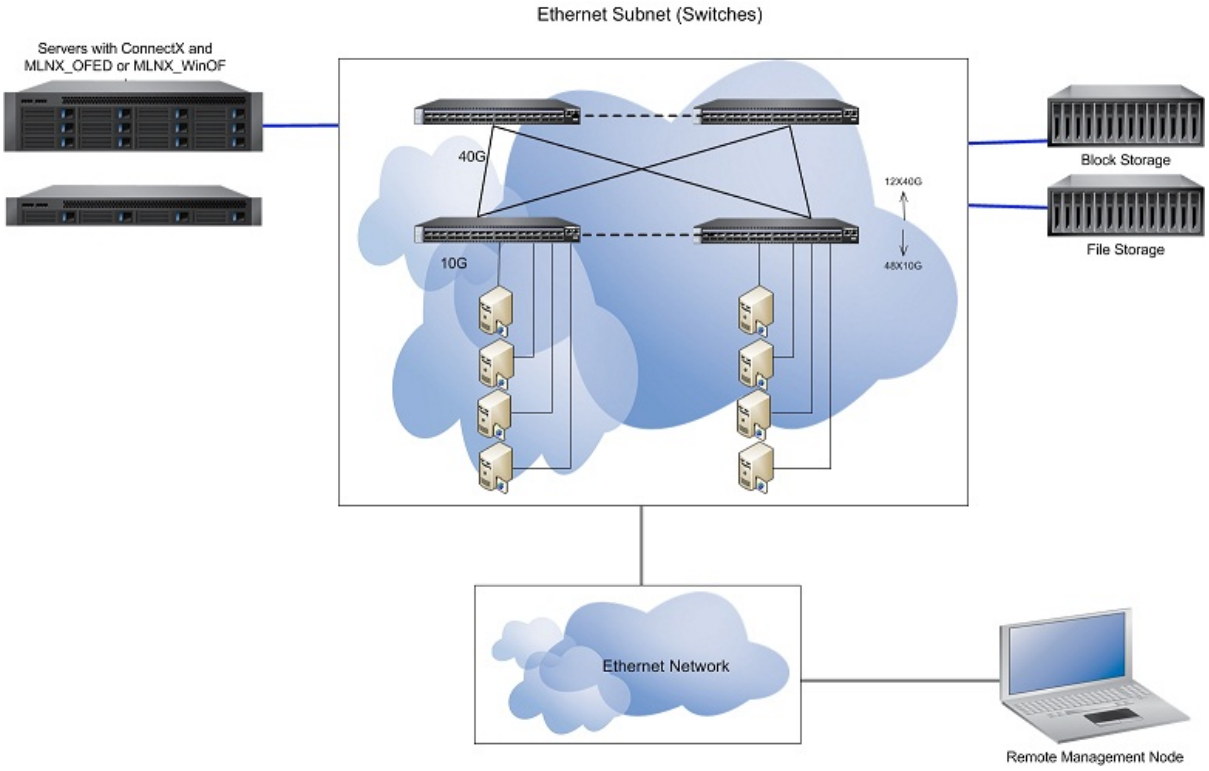

Feature	Detail
Layer 2 Feature Set	• Multi Chassis LAG (MLAG) • IGMP V2/V3, Snooping, Querier • VLAN 802.1Q (4K) • Q-In-Q • 802.1w Rapid Spanning Tree (RSTP) • BPDU Filter, Root Guard • Loop Guard, BPDU Guard • 802.1s Multiple STP (MSTP) • PVRST+ (Rapid Per VLAN STP+) • 802.3ad Link Aggregation (LAG) & LACP • 32 Ports/Channel—64 Groups Per System • Port Isolation • LLDP • Store & Forward / Cut-through mode of work • HLL • 10/25/40/50/100GbE • Jumbo Frames (9216 BYTES) • Unicast MAC addresses

Layer 3 Feature Set	• 64 VRFs • IPv4 & IPv6 Routing inc Route maps: • BGP4, OSPFv2 • PIM-SM & PIM-SSM (inc PIM-SM over MLAG) • BFD (BGP, OSPF, static routes) • VRRP • MAGP • DHCPv4/v6 Relay • Router Port, int Vlan, NULL Interface for Routing • ECMP, 64-way • IGMPv2/v3 Snooping Querier
Synchronization	• PTP IEEE-1588 (SMPTE profile) • NTP
Quality of Service	• 802.3X Flow Control • WRED, Fast ECN & PFC • 802.1Qbb Priority Flow Control • 802.1Qaz ETS • DCBX—App TLV support • Advanced QoS—qualification, rewrite, policers • 802.1AB • Shared buffer management
Management & Automation	• ZTP • Ansible, SALT Stack • FTP \ TFTP \ SCP • AAA , RADIUS \ TACACS+ \ LDAP • JSON & CLI , enhanced web UI • SNMP v1,2,3 • In-band management • DHCP, SSHv2, Telnet • SYSLOG • 10/100/1000 ETH RJ45 MNG ports • USB console port for management • Dual SW image • Events history • ONIE
Network Virtualization	• VXLAN EVPN—L2 stretch use case • VXLAN Hardware VTEP—L2 centralized gateway • Integration with VMware NSX & OpenStack, etc.

Software Defined Network (SDN)	• OpenFlow 1.3:
	  • Hybrid
	  • Supported controllers: ODL, ONOS, FloodLight, RYU, etc.

Docker Container	• Full SDK access through the container • Persistent container & shared storage
Monitoring & Telemetry	• What Just Happened (WJH) • sFlow • Real time queue depth histograms & thresholds • Port mirroring (SPAN & ERSPAN) • Enhanced Link & Phy Monitoring • BER degradation monitor • Enhanced health mechanism • 3rd party integration (Splunk, etc.)

Security	• USA Department of Defense certification—UC APL • System secure mode—FIPS 140-2 compliance • Storm Control • Access Control Lists (ACLs L2-L4 & user defined) • 802.1X—Port Based Network Access Control • SSH server strict mode—NIST 800-181A • CoPP (IP filter) • Port isolation

2 Getting Started

The procedures described in this page assume that you have already installed and powered on your switch according to the instructions in the Hardware Installation Guide, which was shipped with the product.

2.1 Configuring the Switch for the First Time

Due to California Senate Bill No. 327,starting from software version 3.8.2000, Admin and Monitor passwords will need to be typed in manually—no automatic passwords will be created by default.

When the reset button is held for 15 seconds, the management module is reset and the password is deleted. You will then be able to enter without a password and make a new password for the user admin.

Any account created with admin privileges can change all passwords of other user accounts, including other user accounts with admin privileges.

To initialize the switch do the following:

1. Connect the host PC to the console (RJ-45) port of the switch system using the supplied cable.

   DHCP is enabled by default over the MGT port. Therefore, if you have configured your DHCP server and connected an RJ-45 cable to the MGT port, simply log in using the designated IP address.

2. Configure a serial terminal with the settings described below.

   Using NVIDIA cables is manditory.

   This step may be skipped if the DHCP option is used and an IP is already configured for the MGT port.

Parameter	Setting
Baud Rate	115200
Data bits	8
Stop bits	1
Parity	None

Parameter	Setting
Flow Control	None

3. The boot menu is prompted.

```
NVIDIA Onyx Boot Menu:

   1: <image #1>
   2: <image #2>
   u: USB menu (if USB device is connected) (password required)
   c: Command prompt (password required)

   Choice:
```

> Select “0” to boot with software version installed on partition #1.
> Select “1” to boot with software version installed on partition #2.

The boot menu features a countdown timer. It is recommended to allow the timer to run out by not selecting any of the options.

4. Login as admin and use admin as password. If the machine is still initializing, you might not be able to access the CLI until initialization completes. As an indication that initialization is ongoing, a countdown of the number of remaining modules to be configured is displayed in the following format: “<no. of modules> Modules are being configured”.
5. Go through the Switch Management configuration wizard.
   IP configuration by DHCP:

Wizard Session Display (Example)	Comments
`Do you want to use the wizard for initial configuration? yes`	You must perform this configuration the first time you operate the switch or after resetting the switch to the factory defaults. Type “yes” and then press <Enter>.
`Step 1: Hostname? [switch-1]`	If you wish to accept the default hostname, then press <Enter>. Otherwise, type a different hostname and press <Enter>.
`Step 2: Use DHCP on mgmt0 interface? [yes]`	Perform this step to obtain an IP address for the switch. (mgmt0 is the management port of the switch.) - If you wish the DHCP server to assign the IP address, type “yes” and press <Enter>. If you type “no” (no DHCP), then you will be asked whether you wish to use the “zeroconf” configuration or not. If you enter “yes” (yes Zeroconf), the session will continue as shown in the "IP zeroconf configuration" table. If you enter “no” (no Zeroconf), then you need to enter a static IP, and the session will continue as shown in the "Static IP configuration" table.
`Step 3: Enable IPv6 [yes]`	Perform this step to enable IPv6 on management ports. The default is "yes" (enabled). If you enter “no” (no IPv6), then you will automatically be referred to Step 5.

Wizard Session Display (Example)	Comments
`Step 4: Enable IPv6 autoconfig (SLAAC) on mgmt0 interface? [no]`	Perform this step to enable stateless address autoconfig on external management port. The default is "no" (disabled). If you wish to enable it, type “yes” and press <Enter>.
`Step 5: Use DHCPv6 on mgmt0 interface? [yes]`	Perform this step to enable DHCPv6 on the MGMT0 interface.
`Step 6: Update time?`	Perform this step to change the time configured. Press <enter> to leave the current time.
`Step 7: Enable password hardening? [yes]`	Perform this step to enable/disable password hardening on your machine. If enabled, new passwords will be checked upon configured restrictions. The default is "yes" (enabled). If you wish to disable it, enter “no”.
`Step 8: Admin password (Must be typed)? <new_password>`	To avoid illegal access to the machine, please type a password and then press <Enter>. Starting from the 3.8.2000 release, the user must type in the admin password upon initial configuration. Due to Senate Bill No. 327, this stage is required and cannot be skipped.
`Step 9: Confirm admin password? <new_password>`	Confirm the password by re-entering it. Note that password characters are not printed.
`Step 10: Monitor password (Must be typed)? <new_password>`	To avoid illegal access to the machine, please type a password and then press <Enter>. Starting from the 3.8.2000 release, the user must type in the admin password upon initial configuration. Due to Senate Bill No. 327, this stage is required and cannot be skipped.
`Step 11: Confirm monitor password? <new_password>`	Confirm the password by re-entering it. Note that password characters are not printed.

`You have entered the following information:`
`Hostname: <switch name>`
`Use DHCP on mgmt0 interface: yes`
`Enable IPv6: yes`
`Enable IPv6 autoconfig (SLAAC) on mgmt0 interface: yes`
`Enable DHCPv6 on mgmt0 interface: no`
`Update time: <current time>`
`Enable password hardening: yes`
`Admin password (Enter to leave unchanged): (CHANGED)`
`To change an answer, enter the step number to return to.`
`Otherwise hit <enter> to save changes and exit.`
`Choice: <Enter>`
`Configuration changes saved.`
`To return to the wizard from the CLI, enter the “configuration jump-start” command`
`from configuration mode. Launching CLI...`
`<switch name> [standalone: master] >`
	The wizard displays a summary of your choices and then asks you to confirm the choices or to re-edit them.
	Either press <Enter> to save changes and exit, or enter the configuration step number that you wish to return to.
	To run the command “configuration jump-start” you must be in Config mode.

Static IP configuration:

Wizard Session Display (Example)

```
Do you want to use the wizard for initial configuration? y

Step 1: Hostname? [switch-112126]
Step 2: Use DHCP on mgmt0 interface? [yes] n
Step 3: Use zeroconf on mgmt0 interface? [no]
Step 4: Primary IP address? 192.168.10.4
Mask length may not be zero if address is not zero (interface mgmt0)

Step 5: Netmask? [0.0.0.0] 255.255.255.0
Step 6: Default gateway? 192.168.10.1
Step 7: Primary DNS server?
Step 8: Domain name?
Step 9: Enable IPv6? [yes] yes
Step 10: Enable IPv6 autoconfig (SLAAC) on mgmt0 interface? [no] no
Step 11: Update time? [yyyy/mm/dd hh:mm:ss]
Step 12: Enable password hardening? [yes] yes
Step 13: Admin password (Enter to leave unchanged)?

You have entered the following information:

Hostname: switch-112126
Use DHCP on mgmt0 interface: no
Use zeroconf on mgmt0 interface: no
Primary IP address: 192.168.10.4
Netmask: 255.255.255.0
Default gateway: 192.168.10.1
Primary DNS server:
Domain name:
Enable IPv6: yes
Enable IPv6 autoconfig (SLAAC) on mgmt0 interface: no
Update time: yyyy/mm/dd hh:mm:ss
Enable password hardening: yes
Admin password (Enter to leave unchanged): (unchanged)

To change an answer, enter the step number to return to.
Otherwise hit <enter> to save changes and exit.

Choice:
Configuration changes saved.
To return to the wizard from the CLI, enter the “configuration jump-start” command from configure
mode. Launching CLI...
<hostname>[standalone: master] >
```

IP zeroconf configuration:

Wizard Session Display (Example)

```
Configuration wizard

Do you want to use the wizard for initial configuration? y

Step 1: Hostname? [switch-112126]
Step 2: Use DHCP on mgmt0 interface? [no]
Step 3: Use zeroconf on mgmt0 interface? [no] yes
Step 4: Default gateway? [192.168.10.1]
Step 5: Primary DNS server?
Step 6: Domain name?
Step 7: Enable IPv6? [yes] yes
Step 8: Enable IPv6 autoconfig (SLAAC) on mgmt0 interface? [no] no
Step 9: Update time? [yyyy/mm/dd hh:mm:ss]
Step 10: Admin password (Enter to leave unchanged)?

You have entered the following information:

Hostname: switch-112126
Use DHCP on mgmt0 interface: no
Use zeroconf on mgmt0 interface: yes
Default gateway: 192.168.10.1
Primary DNS server:
Domain name:
Enable IPv6: yes
Enable IPv6 autoconfig (SLAAC) on mgmt0 interface: yes
Update time: yyyy/mm/dd hh:mm:ss
Enable password hardening: yes
Admin password (Enter to leave unchanged): (unchanged)

To change an answer, enter the step number to return to.
Otherwise hit <enter> to save changes and exit.

Choice:

Configuration changes saved.

To return to the wizard from the CLI, enter the "configuration jump-start"
command from configure mode. Launching CLI...
<hostname> [standalone: master] >
```

6. Check the mgmt0 interface configuration before attempting a remote (for example, SSH) connection to the switch. Specifically, verify the existence of an IP address.

```
switch # show interfaces mgmt0

Interface mgmt0 status:
  Comment          :
  Admin up         : yes
  Link up          : yes
  DHCP running     : yes
  IP address       : 10.12.67.34
  Netmask          : 255.255.0.0
  IPv6 enabled     : yes
  Autoconf enabled: no
  Autoconf route   : yes
  Autoconf privacy: no
  DHCPv6 running   : no
  IPv6 addresses   : 1

  IPv6 address:
    fe80::268a:7ff:fe53:3d8e/64

  Speed            : 1000Mb/s (auto)
  Duplex           : full (auto)
  Interface type   : ethernet
  Interface source: physical
  MTU              : 1500
  HW address       : 00:02:c9:11:a1:b2

  Rx:
   11700449  bytes
      55753  packets
          0  mcast packets
          0  discards
          0  errors
```

```
          0  overruns
          0  frame

    Tx:
    5139846  bytes
      28452  packets
          0  discards
          0  errors
          0  overruns
          0  carrier
          0  collisions
       1000  queue len
```

2.1.1 Configuring the Switch with ZTP

Zero-touch Provisioning (ZTP) automates initial configuration of switch systems at boot time. It helps minimize manual operation and reduce customer initial deployment cost.

For more information, please refer to section "Zero-touch Provisioning".

2.1.2 Rerunning the Wizard

To rerun the wizard:

1. Enter Config mode. Run:

```
switch > enable
switch # config terminal
```

2. Rerun the wizard. Run:

```
switch (config) # configuration jump-start
```

2.2 Starting the Command Line (CLI)

1. Set up an Ethernet connection between the switch and a local network machine using a standard RJ-45 connector.
2. Start a remote secured shell (SSH) to the switch using the command "ssh -l <username> <switch ip address>".

```
rem_mach1 >  ssh -l <username> <ip address>
```

3. Log into the switch (default username is admin, password admin).
4. Read and accept the EULA when prompted.
5. Once the following prompt appears, the system is ready to use.

```
NVIDIA Onyx Switch Management

Password:
Last login: <time> from <ip-address>


NVIDIA Switch
Please read and accept the End User License Agreement located at:
https://www.mellanox.com/related-docs/prod_management_software/MLNX_Onyx_EULA.pdf
switch >
```

2.3 Starting the Web User Interface (WebUI)

To start a WebUI connection to the switch platform, follow the steps below:

> WebUI access is enabled by default. To disable web access, run the command “no web http enable” or “no web https enable” on the CLI.

1. Set up an Ethernet connection between the switch and a local network machine using a standard RJ-45 connector.
2. Open a web browser that is Firefox, Chrome, Internet Explorer, or Safari.

   > Make sure the screen resolution is set to 1024*768 or higher.

   > In order to access WebUI through Sarafi 5.3, enable http:

   ```
   no web https ssl secure-cookie enable
   web http enable
   ```

3. Type the IP address of the switch or its DNS name in the following format: https://<switch_IP_address>.
4. Log into the switch (default user name is admin, password admin).

5. Read and accept the EULA, if prompted.
   The prompt will only occur if the switch has never been accessed through the CLI before.

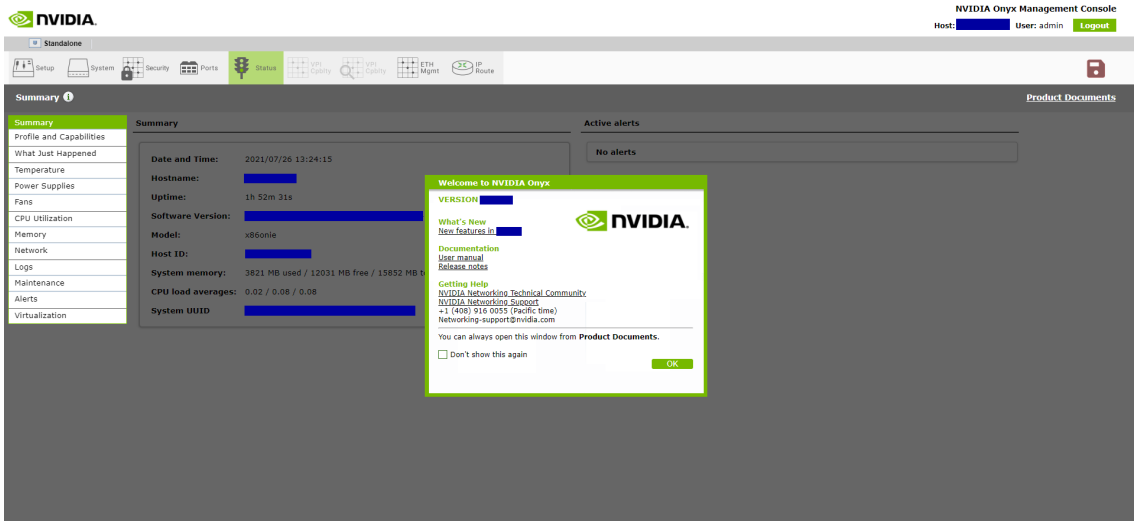

6. The Welcome popup appears. After reading through the content, click OK to continue.
   To reach the OS documentation, click on the links under the Documentation heading.
   The link under What's New takes leads to the Changes and New Features section of the switch OS Release Notes. You may also tick the box to not show this popup again. To see this window again, click "Product Documents" on the upper right corner of the WebUI.

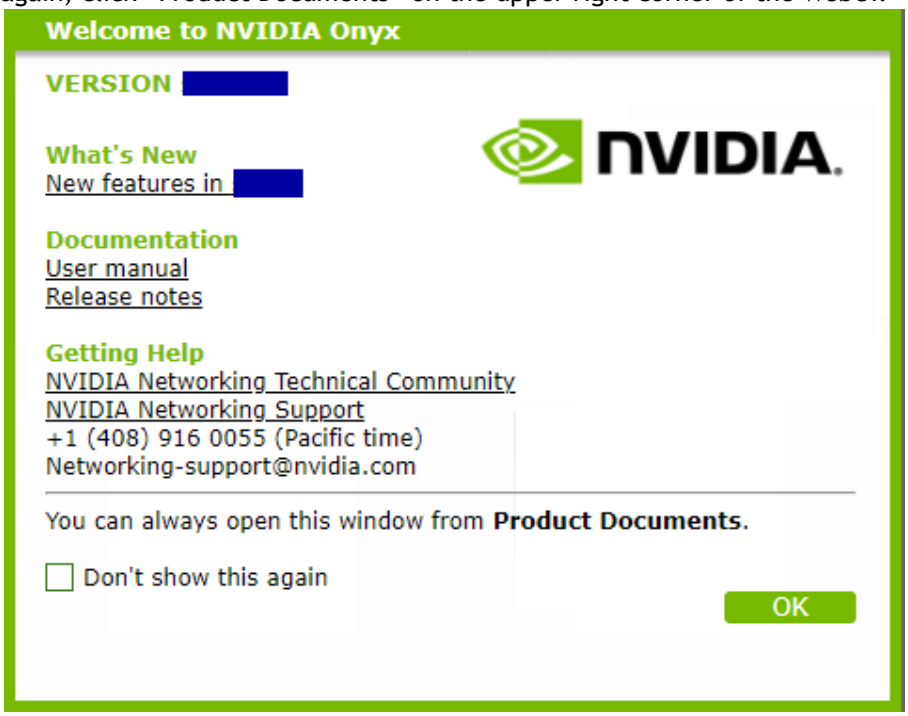

7. A default status summary is displayed.

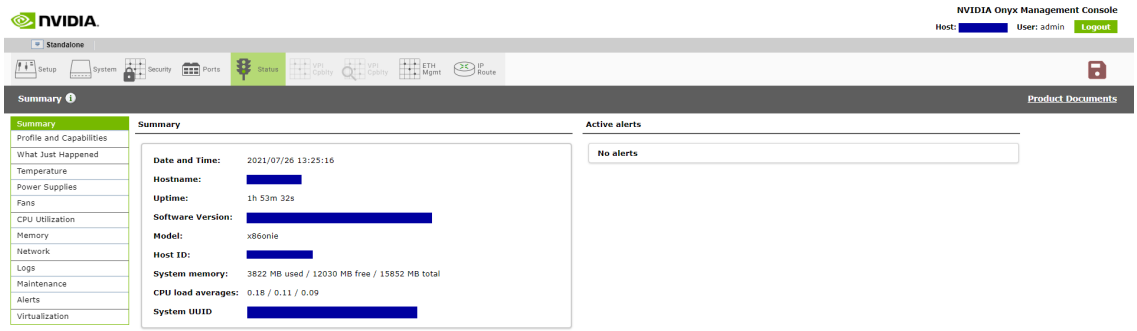

2.4 Zero-touch Provisioning

Zero-Touch Provisioning (ZTP) automates initial configuration of switch systems at boot time. It helps minimize manual operation and reduce customer initial deployment cost. ZTP allows for automatic upgrade of the switch with a specified OS image, setting up initial configuration database, and to load and run a container image file.

The initial configuration is applied using a regular text file. The user can create such a configuration file by editing the output of a “show running-config” command.

> Only a textual configuration file is supported.

The user-defined docker image can be used by customers to run their own applications in a sandbox on their platform. They can therefore also be used for automating initial configuration.

> Only one docker container can be launched in ZTP.

2.4.1 Running DHCP-ZTP

There is no explicit command to enable ZTP. It is enabled by default. Disabling it is performed by a user-initiated configuration save (using the command “configuration write”). The only way to re-enable ZTP is to run a “reset factory” command, clearing the configuration of the switch and rebooting the system.

ZTP is based on DHCP. For ZTP to work, the software enables DHCP by default on all its management interfaces. The switch OS requests option 66 (tftp-server-name) and 67 (bootfile-name) from the DHCPv4 server or option 58 (bootfile-url) from the DHCPv6 server, and waits for the DHCP responses containing file URLs. The DHCP server must be configured to send back the URLs for the software image, configuration file, and docker container image via these two options. Option 66 would contain the URL prefix to the location of the files, option 67 would contain the name of files, and option 58 would contain the complete URLs of files. The format of these two options is a string list separated by commas. The list items are placed in a fixed order:

DHCPv4

```
option tftp-server-name "<image server url>, <config server url>, <docker container server url>";
option bootfile-name "<image file>, <config file>, <docker container file>";
```

DHCPv6

```
option dhcp6.bootfile-url "<image server url/image file>, <config server url/config file>, <docker container server
url/docher container file>";
```

> The item value can be empty, but the comma shall not be omitted.

The item value can be empty, but the comma shall not be omitted.

To have DHCP server discern the proper files based on switch-specific information, the OS must provide identifying information for the server to classify the switches. In addition, the OS attaches option 43 (vendor-specific information) and option 60 (vendor class identifier) in DHCPv4 requests and option 17 (vendor-opts) in DHCPv6. Option 60 is set as string “Mellanox” and options 17 and 43 contain the following specific sub-options:

- System Model
- Chassis Part Number
- Chassis Serial Number
- Management MAC
- System Profile
- NVIDIA Onyx Release Version

The corresponding subtypes respectively are defined as:

```
DHCP_VENDOR_ENCAPSULATED_SUBOPTION_TLV_TYPE_MODEL       1
DHCP_VENDOR_ENCAPSULATED_SUBOPTION_TLV_TYPE_PARTNUM     2
DHCP_VENDOR_ENCAPSULATED_SUBOPTION_TLV_TYPE_SERIAL      3
DHCP_VENDOR_ENCAPSULATED_SUBOPTION_TLV_TYPE_MAC         4
DHCP_VENDOR_ENCAPSULATED_SUBOPTION_TLV_TYPE_PROFILE     5
DHCP_VENDOR_ENCAPSULATED_SUBOPTION_TLV_TYPE_RELEASE     6
```

Upon receiving such DHCP requests from a client, the server should be able to map the switch-specific information to the target file URLs according to predefined rules.

Once the OS receives the URLs from the DHCP server, it executes ZTP as follows:

If the software image URL is not specified, this step is skipped. Otherwise:

a. Perform disk space cleanup if necessary and fetch the image if it does not exist locally
b. Resolve the image version:
c. If it is already installed on active partition, proceed to step 2
d. If it is installed on a standby partition, switch partition and reboot
e. If it is not installed locally, install it and switch to the new image and then reboot
f. If a reboot occurs, ZTP performs step 1 again and no image upgrade will occur

If configuration file URL is not specified, skip this step. Otherwise:

a. Fetch the configuration file
b. Apply the configuration file

Skip these steps if a docker image file URL is not specified. Otherwise:

a. Fetch the docker image file
b. Load the docker image
c. Clean up the docker images with the same name and different tag.
d. Start the container based on the image
e. Remove the downloaded docker image file

While performing file transfer via HTTP, the same information as DHCP option 43 is expected to be carried in a HTTP GET request. This switch software supports the following proprietary HTTP headers:

- MlnxSysProfile
- MlnxMgmtMac
- MlnxSerialNumber
- MlnxModelName
- MlnxPartNumber
- MlnxReleaseVersion

If some sort of failure occurs, the switch waits a random number of seconds between 1 and 20 and reattempts the operation. The switch attempts this up to 10 times.
ZTP progress is printed to terminals including console and active SSH sessions.

2.4.2 ZTP and OS Upgrade

Software upgrade from non-ZTP versions to ZTP versions and vice versa is supported. When upgrading from a non-ZTP version, ZTP is disabled because ZTP is always assumed to start with an empty configuration, otherwise the final configuration becomes a mixture of the existing configuration from the stored database and new configuration from the server and hence not deterministic.

2.4.3 DHCPv4 Configuration Example

The following is a URL configuration example for ISC DHCPv4 server:

```
host master {
        hardware ethernet E4:1D:2D:5B:72:80;
        fixed-address 3.1.2.13;
        option tftp-server-name "scp://<user>:<password>@3.1.3.100/ztp/,scp://
                                 <user>:<password>@3.1.3.100/ztp/,scp://
                                 <user>:<password>@3.1.3.100/ztp/";
        option bootfile-name "image-X86_64-3.6.4612.img, switch-1.conf, ubuntu.img.gz";
   }
```

DHCPv4 request is made out of the following components:

- Option 43 (vendor-encapsulated-options) and option 60 (vendor-class-identifier) are added in the DHCPv4 request packet
- Option 66 (tftp-server-name) and option 67 (bootfile-name) are added in the parameter request list of DHCPv4 request packet

2.4.4 DHCPv6 Configuration Example

The following is a DHCPv6 configuration example:

```
host master {
        ......
       option dhcp6.bootfile-url "scp://<user>:<password>@[2000::1]/ztp/image-X86_64-
                                 3.6.4612.img, scp://<user>:<password>@[2000::1]/ztp/
                                 switch.conf, scp://<user>:<password>@[2000::1]/ztp/
                                 ubuntu.img.gz";
   }
```

```
host master {
        ......
       option dhcp6.bootfile-url "scp://<user>:<password>@[2000::1]/ztp/image-X86_64-
                                 23.01.0100.img, scp://<user>:<password>@[2000::1]/ztp/
                                 switch.conf, scp://<user>:<password>@[2000::1]/ztp/
                                 ubuntu.img.gz";
   }
```

DHCPv6 request is made out of the following components:

- Option 17 (vendor-opts) is added in the DHCPv6 request packet
- Option 59 (bootfile-url) is added in the parameter request list of DHCPv6 request packet

2.4.5 ZTP Commands

2.4.5.1 no zero-touch suppress-write

	no zero-touch suppress-write Disables suppression of configuration write.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config
History	3.6.5000 3.9.2400: Added note
Example	switch (config) # no zero-touch suppress-write
Related Commands	show zero-touch

Notes	• When ZTP is active, “configuration write” is suppressed because it may interfere with ZTP operation. Therefore, after running “no zero-touch suppress-write” if “configuration write” is performed, then ZTP is disabled as a consequence of the database save.
	• To automatically save the configuration at the end of applying a configuration via ZTP, append the following two commands to the end of the config files. The first command will turn off the ZTP suppress-write, then the configuration write command should work.
	  ◦ `no zero-touch suppress-write`
	  ◦ `configuration write`

2.4.5.2 zero-touch abort

	zero-touch abort Aborts on-going zero-touch process.
Syntax Description	N/A
Default	Enabled

Configuration Mode	config
History	3.6.5000
Example	switch (config) # zero-touch abort Zero-touch failed [Zero-touch is aborted by operator] Zero-touch provisioning will be aborted
Related Commands	show zero-touch
Notes	

2.4.5.3 show zero-touch

	show zero-touch Displays zero-touch status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.5000
Example	switch (config) # show zero-touch Zero-Touch status: Active: yes Status: Waiting for zero-touch start Suppress-write: no Configured by zero-touch: no Configuration changed after zero-touch: no
Related Commands	zero-touch abort zero-touch suppress-write
Notes	

3 User Interfaces

The following pages provide information on the interfaces available for to manage and validate the status of the system.

- LED Indicators
- Command Line Interface (CLI)
- Secure Shell (SSH)
- Web Interface Overview
- UI Commands

3.1 LED Indicators

For information regarding LED indicators, go to the link of the relevant ASIC:

- SN2000 system LED indicators
- SN3000 system LED indicators
- SN4000 system LED indicators

3.2 Command Line Interface (CLI)

NVIDIA Onyx is equipped with an industry-standard command line interface (CLI). The CLI is accessed through SSH or Telnet sessions or directly through the console port on the front panel, if it exists.

3.2.1 CLI Modes

The CLI can be in one of various modes. Each of the modes makes available a certain group (or level) of commands for execution. The following are some of the CLI configuration modes:

Configuration Mode	Description
Standard	When the CLI is launched, it begins in Standard mode. This is the most restrictive mode and only has commands to query a restricted set of state information. Users cannot take any actions that directly affect the system, nor can they change any configuration.
Enable	The "enable" command moves the user to Enable mode. This mode offers commands to view all state information and take actions like rebooting the system, but it does not allow any configurations to be changed. Its commands are a superset of those in Standard mode.
config	The "configure terminal" command moves the user from Enable mode to Config mode. Config mode is allowed only for user accounts with the “admin” role (or capabilities). This mode has a full unrestricted set of commands to view anything, take any action, and change any configuration. Its commands are a superset of those in Enable mode. To return to Enable mode, enter the command "exit" or "no configure". Note that moving directly between Standard and Config mode is not possible.

Configuration Mode	Description
config interface management	Configuration mode for management interface mgmt0, mgmt1, and loopback
config interface ethernet	Configuration mode for Ethernet interface
config interface port-channel	Configuration mode for Port channel (LAG)
config vlan	Configuration mode for VLAN
Any command mode	Several commands, such as “show” commands, can be applied within any context
"no" parameter	When the "no" form of the command is used, the command is erased from the running-config and reverts to either the default or inherited value. Note that if used on a string (e.g., password), that value is either removed unless it can be inherited. If used on a boolean value, it is FALSE unless it has either a default or an inherited value. See example in "Using the “no” Command Form" section.
“disable” parameter	When the "disable" form of the command is used, it creates an entry in running-config that prevents inheritance and reverts to the default system settings. If used on a string (e.g., password), that value is removed (it cannot be inherited). If used on a boolean value, the value is set to FALSE (it cannot be inherited).

3.2.2 Syntax Conventions

To help identify the different parts of a CLI command, the following table explains conventions of presenting the syntax of commands.

Syntax Convention	Description	Example
< > Angled brackets	Indicate a value/variable that must be replaced.	<1...65535> or <interface>
[ ] Square brackets	Indicate optional parameters. Only one parameter out of the parameters listed with in the brackets can be used—the user cannot have a combination of the parameters unless stated otherwise.	[destination-ip | destination-port | destination-mac]
{ } Braces	Indicate alternatives or variables that are required for the parameter in square brackets.	[mode {active | on | passive}]
| Vertical bars	Identify mutually exclusive choices.	active | on | passive

Do not use the angled or square brackets, vertical bar, or braces in command lines. This guide uses these symbols only to show the different entry types.

CLI commands and options are in lowercase and are case-sensitive.
For example, when entering the enable command, "enable" must be all in lowercase; it cannot be ENABLE or Enable. Text entries created are also case-sensitive.

3.2.3 Getting Help

Context-sensitive help may be requested at any time by pressing “?” in the command line. This will show a list of choices for the word that is currently selected or, if nothing has been typed yet, will show a list of top-level commands.
For example, typing "?" in the command line in Standard mode, will provide a link of the following available commands.

```
switch > ?
cli             Configure CLI shell options
enable          Enter enable mode
exit            Log out of the CLI
help            View description of the interactive help system
no              Negate or clear certain configuration options
show            Display system configuration or statistics
slogin          Log into another system securely using ssh
switch          Configure switch on system
telnet          Log into another system using telnet
terminal        Set terminal parameters
traceroute      Trace the route packets take to a destination
switch >
```

Typing a legal string and then pressing “?” without a space character before it, will provide either a description of the command that was typed so far or the possible command/parameter completions. Typing “?” after a space character and “<cr>” is shown, means that, so far, a complete command has been typed. Pressing Enter (carriage return) will execute the command.

Try the following, to get started:

```
?
show ?
show c?
show clock?
show clock ?
show interfaces ?     (from enable mode)
```

Enter “help” to view a description of the interactive help system.
Note also that the CLI supports command and/or parameter tab-completions and their shortened forms. For example, you can enter “en” instead of the “enable” command, or “cli cl” instead of “cli clear-history”. In case of ambiguity (in case more than one completion option is available), press Tabs twice to obtain the disambiguation options. Thus, to learn which commands start with the letter “c”, type “c” and click twice on the Tab key to get the following:

```
switch # c<tab>
clear       cli       configure
switch # c
```

This signifies that there are three commands that start with the letter “c”: "clear", "cli", and "configure".

3.2.4 Prompt and Response Conventions

The prompt always begins with the hostname of the system. What follows depends on what command mode the user is in. To demonstrate by example, assuming the machine name is “switch”, the prompts for each of the modes are:

```
switch >              (Standard mode)
switch #              (Enable mode)
switch (config) #     (Config mode)
```

The following session shows how to move between command modes:

```
switch >                        (You start in Standard mode)
switch > enable                 (Move to Enable mode)
switch #                        (You are in Enable mode)
switch # configure terminal     (Move to Config mode)
switch (config) #               (You are in Config mode)
switch (config) # exit          (Exit Config mode)
switch #                        (You are back in Enable mode)
switch # disable                (Exit Enable mode)
switch >                        (You are back in Standard mode)
```

Commands entered do not print any response and simply show the command prompt after pressing <Enter>.

If an error is encountered while executing a command, the response will begin with “%”, followed by a description of the error.

3.2.5 Using the “no” Command Form

Several Config commands use the “no” form of the command to reset a parameter value to its inherited, or default, value.

The command sequence below performs the following:

1. Displays the current CLI session option.
2. Disables auto-logout.
3. Displays the new CLI session options (auto-logout is disabled).
4. Re-enables auto-logout (after 15 minutes).
5. Displays the final CLI session options (auto-logout is enabled).

```
// 1. Display the current CLI session options
switch (config) # show cli
CLI current session settings:
Maximum line size: 8192
Terminal width: 157 columns
Terminal length: 60 rows
Terminal type: xterm
Auto-logout: 15 minutes
Paging: enabled
Progress tracking: enabled
Prefix modes: enabled
...
// 2. Disable auto-logout
switch (config) # no cli session auto-logout
// 3. Display the new CLI session options
switch (config) # show cli
CLI current session settings:
Maximum line size: 8192
Terminal width: 157 columns
Terminal length: 60 rows
Terminal type: xterm
Auto-logout: disabled
Paging: enabled
Progress tracking: enabled
Prefix modes: enabled
...
// 4. Re-enable auto-logout after 15 minutes
switch (config) # cli session auto-logout 15
// 5. Display the final CLI session options
switch (config) # show cli
CLI current session settings:
Maximum line size: 8192
Terminal width: 157 columns
Terminal length: 60 rows
Terminal type: xterm
Auto-logout: 15 minutes
Paging: enabled
Progress tracking: enabled
Prefix modes: enabled
...
```

3.2.6 Parameter Key

This page provides a key to the meaning and format of angle-bracketed parameters in the commands that are listed in this document.

Parameter	Description
<domain>	A domain name
<hostname>	A hostname (e.g., “switch-1”)
<ifname>	An interface name (e.g., “mgmt0”, “mgmt1”, “lo” (loopback), and so forth).
<index>	A number to be associated with aliased (secondary) IP addresses.
<IP address>	An IPv4 address (e.g., "192.168.0.1")
<log level>	A syslog logging severity level. Possible values, from least to most severe, are as follows: “debug”, “info”, “notice”, “warning”, “error”, “crit”, “alert”, “emerg”.
<GUID>	Globally unique identifier. A number that uniquely identifies a device or component.
<MAC address>	A MAC address. The segments may be 8 bits or 16 bits at a time, and may be delimited by “:” or “.” (e.g., “11:22:33:44:55:66”, “1122:3344:5566”, “11.22.33.44.55.66”, or “1122.3344.5566”).
<netmask>	A netmask (e.g., “255.255.255.0”) or mask length prefixed with a slash (e.g., “/24”). Both examples express the same information in different formats.
<network prefix>	An IPv4 network prefix specifying a network. Used in conjunction with a netmask to determine which bits are significant. e.g., “192.168.0.0”.
<regular expression>	An extended regular expression as defined by the “grep” in the main page. (The value provided here is passed on to “grep -E”.)
<node id>	ID of a node belonging to a cluster. This is a numerical value greater than zero.
<cluster id>	A string specifying the name of a cluster.
<port>	TCP/UDP port number.
<TCP port>	A TCP port number in the full allowable range [0...65535].
<URL>	A normal URL, using any protocol that wget supports, including HTTP, HTTPS, FTP, SFTP, and TFTP or a pseudo-URL specifying an scp file transfer. The scp pseudo-URL format is scp://username:password@hostname/path/filename. Note that the path is an absolute path. Paths relative to the user's home directory are not currently supported. Because the implementation of FTP does not support authentication, use SCP or SFTP for that. Note also that omitting “:password” part, may require entering the password in a follow-up prompt, where it can be typed in securely (without the characters being echoed). This prompt will occur if the “cli default prompt empty-password” setting is true; otherwise, the CLI will assume that no password is desired. Including the “:” character, will be taken as an explicit declaration that the password is empty and no prompt will appear.

3.2.7 CLI Pipeline Operator Commands

3.2.7.1 CLI Filtration Options “include” and “exclude”

The NVIDIA Onyx CLI supports filtering “show” commands to display lines containing or excluding certain phrases or characters. To filter the outputs of the “show” commands use the following format:

```
switch (config) # <show command> | {include | exclude} <extended regular expression> [<ignore-case>] [next <lines>]
[prev <lines>]
```

The filtering parameters are separated from the show command they filter by a pipe character (“|”). Quotation marks may be used to include or exclude a string including space, and multiple filters can be used simultaneously as shown in the example below.

```
switch (config) # <show command> | {include <extended regular expression>} [<ignore-case>] [next <lines>] [prev
<lines>] | exclude <extended regular expression> [<ignore-case>] [next <lines>] [prev <lines>]]
```

Example:

```
switch (config) # show asic-version | include SPC
MGMT               SPC                13.1601.3150

switch (config) # show module | exclude PS
======================
 Module    Status
======================
 MGMT      ready
 FAN1      ready
 FAN2      ready

switch (config) # show interfaces | include "Eth|discard pac"
Eth1/1
0 discard packets
0 discard packets
Eth1/2
0 discard packets
0 discard packets
Eth1/3
0 discard packets
0 discard packets
Eth1/4
0 discard packets
0 discard packets
switch (config) # show interfaces | include "Tx" next 5 | exclude broad
Tx
0 packets
0 unicast packets
0 multicast packets
0 bytes
--
```

3.2.7.2 CLI Monitoring Option “watch”

NVIDIA Onyx

```
switch (config) # <show command> | watch [diff] [interval <1-100 secs>]
```

Running this command displays a show-command output that is updated at a time interval specified by the “interval” parameter (2 seconds is the default).

The “diff” parameter highlights the differences between each iteration of the command.

For example running the command “show power | watch diff interval 1” yields something similar to the following:

```
-------------------------------------------------------------------------
Module  Device          Sensor  Power   Voltage  Current  Feed  Status
                                [Watts] [Watts]  [Amp]
-------------------------------------------------------------------------
PS1     power-mon       input   85.00   230.00   0.38     AC    OK
PS2     power-mon       -       -       -        -        -     FAIL

Total power used : 85.00 Watts
Total power capacity : 460.00 Watts
Total power available : 375.00 Watts
Maximum consumed power of all turned on modules: 46.00 Watts
```

With the highlighted black blocks indicating the change that has occurred between one iteration of the command from one second to the next.

To exit “watch” mode, press Ctrl+C.
The “watch” option may be used in conjunction with the “include” and “exclude” options as follows:

```
switch (config) # <show command> | {include | exclude} <extended regular expression> | watch [diff] [interval <1-10
0 secs>]
```

Example:

```
switch (config) # show power | include PS | watch diff interval 1
```

It is possible to count the number of lines in an output of a “show” command by using the following command:

```
switch (config) # <show command> | count
```

Example:

```
switch (config) # show clock
Time:       16:05:43
Date:       2020/05/25
Time zone:  UTC (Etc/UTC)
UTC offset: same as UTC
# show clock | count
4
```

3.2.7.3 CLI “json-print” Option

The NVIDIA Onyx CLI supports printing “show” commands in JSON syntax.
To print the output of the “show” commands as JSON, use the following format:

```
switch (config) # <show command> | json-print
```

Running the command displays an output of the “show” command in JSON syntax structure instead of its regular format. See the following as an example:

```
switch (config) # show system profile
Profile: eth-single-switch
switch (config) # show system profile | json-print
{
"Profile": "eth-single-switch"
}
```

The “json-print” option cannot be used together with filtering (“include” and “exclude”) and/or monitoring (“watch”).

For more information on JSON usage, please refer to “JSON API”.

3.2.7.4 CLI Shortcuts

The following table presents the available keyboard shortcuts on the NVIDIA Onyx CLI.

Key Combination	Description
Ctrl-a	Move cursor to beginning of line
Ctrl-b	Move cursor backward one character without deleting
Ctrl-c	Terminate operation
Ctrl-d	If cursor is in the middle of the line, delete one character forward If cursor is at the end of the line, show autocomplete options for current word or word fragment If cursor at an empty line, same as Esc
Ctrl-e	Move cursor to end of line
Ctrl-f	Move cursor forward one character
Ctrl-h	Delete one character backwards from cursor
Ctrl-i	Auto-complete current word (same as TAB)
Ctrl-j	Return carriage (same as ENTER)
Ctrl-k	Delete line after cursor
Ctrl-l	Clear screen and show line at the top of terminal window
Ctrl-m	Return carriage (same as ENTER)
Ctrl-n	Next line (same as DOWN ARROW)
Ctrl-p	Next line (same as UP ARROW)
Ctrl-t	Transpose the two characters on either side of cursor
Ctrl-u	Delete line
Ctrl-w	Delete the last word
Ctrl-y	Retrieve (“yank”) last item deleted
Esc b	Move cursor one word backward
Esc c	Capitalizes first letter in word after cursor
Esc d	Delete one word forward from cursor
Esc f	Move one word forward from cursor
Esc l	Change word after cursor to lowercase letters
Esc Ctrl-h	Delete one word backward from cursor

Key Combination	Description
Esc [A	Next line (same as DOWN ARROW)
Esc [B	Next line (same as UP ARROW)
Esc [C	Move forward one character from cursor
Esc [D	Move backward one character from cursor

3.3 Secure Shell (SSH)

It is recommended not to use more than 50 concurrent SSH sessions.

3.3.1 Adding a Host and Providing an SSH Key

To add entries to the global known-hosts configuration file and its SSH value, do the following.

1. Change to Config mode.

```
switch > enable
switch # configure terminal
switch (config) #
```

2. Add an entry to the global known-hosts configuration file and its SSH value.

```
switch (config) # ssh client global known-host "myserver ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAIEAsXeklqc8T0EN2mnMcVcfhueaRYzIVqt4rVsrERIjmlJh4mkYYIa8hGGikNa+t5xw2dRrNxnHYLK51bU
sSG1ZNwZT1Dpme3pAZeMY7G4ZMgGIW9xOuaXgAA3eBeoUjFdi6+1BqchWk0nTb+gMfI/MK/heQNns7AtTrvqg/O5ryIc="
```

3. Verify what keys exist in the host.

```
switch (config) # show ssh client
SSH client Strict Hostkey Checking: ask

SSH Global Known Hosts:
    Entry 1: myserver
          Finger Print: d5:d7:be:d7:6c:b1:e4:16:df:61:25:2f:b1:53:a1:06

No SSH user identities configured.

No SSH authorized keys configured.
```

RSA2 and a DSA2 host keys are generated by default. The RSA2 key can be used as SSH server and client, while DSA2 key can only be used as SSH client.
When the switch is a server, use RSA key to connect to the device.

When the switch is a client (e.g., downloading image or uploading logs), RSA key is recommended. DSA key is only for legacy devices and has been deprecated by OpenSSH starting with the 7.0 release.

3.3.2 Retrieving Return Codes When Executing Remote Commands

To stop the CLI and set the system to send return errors if some commands fail, do the following.

1. Connect to the system from the host SSH.
2. Add the flag "-h" after "cli" to notify the system to halt on failure and pass through the exit code.

```
ssh <username>@<hostname> cli -h '"enable" "show interfaces brief"'
```

3.4 Web Interface Overview

The NVIDIA Onyx package equipped with web-based GUI that accepts input and provides output by generating webpages that can be viewed by the user using a web browser.

> The maximum allowed number of WebUI session is 225. Trying to open new sessions beyond this limitation is rejected.

3.4.1 Password Hardening

Upon initial login through the web interface, if the initial login was not completed through the CLI the following prompt will appear (by default, password hardening is enabled).

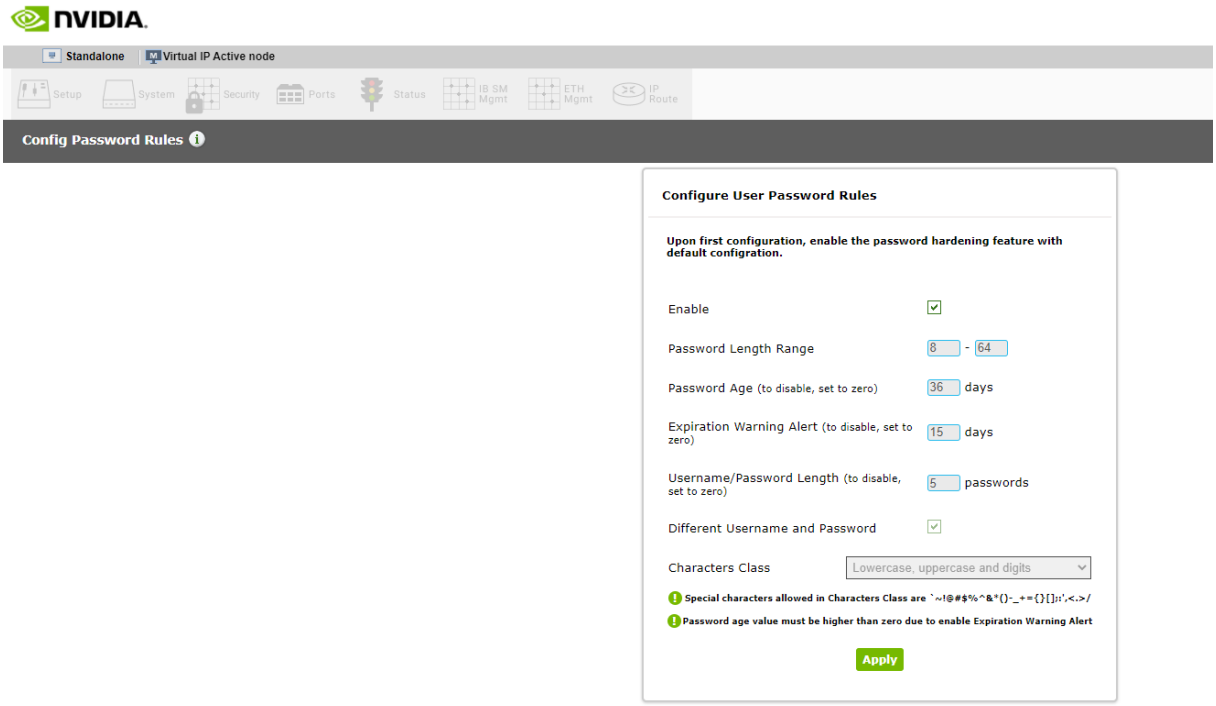

3.4.2 Changing Default Password

The password may be required to be changed upon initial login through the web interface if initial login was not completed through the CLI.

Upon initial login do the following:

1. Login as admin.
2. If the following screen appears (this screen will appear if default password was never changed), type in a new password ("admin" may be reused as the new password).

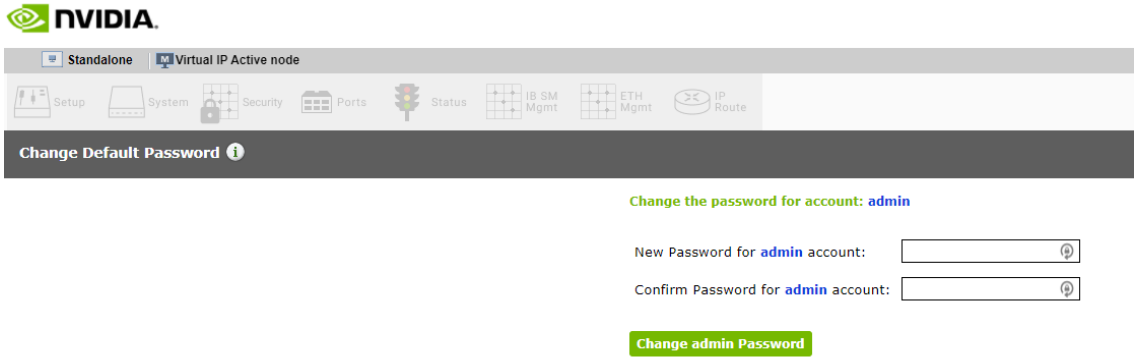

3. Only after successfully changing the admin password (this must be done first), change the monitor password. If the password is not changed, all pages (besides the logout page) will be

locked.

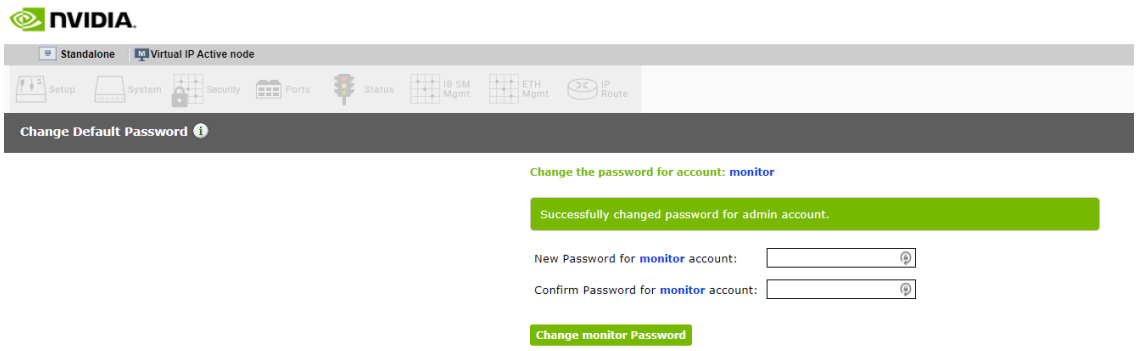

4. After successfully changing the monitor password, the home page may be accessed and the system may be used.

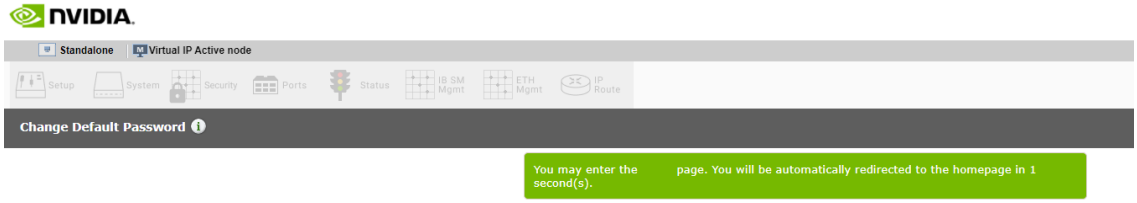

5. Click on the home page link or wait 5 seconds until the countdown reaches 0 and the page is redirected automatically.

> Warning: Entering the monitor user before the default password is changed will block the system (all pages besides the logout page will be blocked).

3.4.3 About Web UI

The web interface makes available the following perspective tabs:

- Setup
- System
- Security
- Ports
- Status
- Ethernet Management
- IP Route

> Make sure to save your changes before switching between menus or submenus. Click the “Save” button to the right of “Save Changes?”.

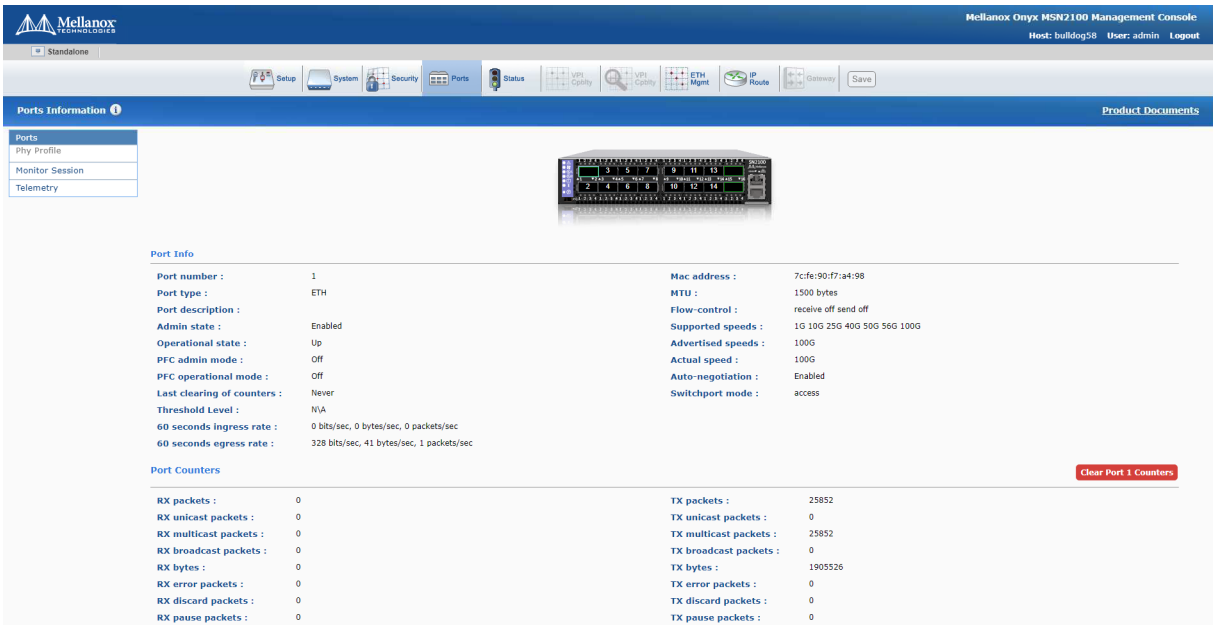

3.4.4 Setup Menu

The Setup menu makes available the following submenus (listed in order of appearance from top to bottom):

Submenu Title	Description
Interfaces	Obtains the status of, configures, or disables interfaces to the fabric. Thus, you can: set or clear the IP address and netmask of an interface; enable DHCP to dynamically assign the IP address and netmask; and set interface attributes such as MTU, speed, duplex, etc.
Routing	Configures, removes or displays the default gateway, and the static and dynamic routes
Hostname	Configures or modifies the hostname Configures or deletes static hosts Note: Changing hostname stamps a new HTTPS certificate
DNS	Configures, removes, modifies or displays static and dynamic name servers
Login Messages	Edits the login messages: Message of the Day (MOTD), Remote Login message, and Local Login message
Address Resolution	Adds static and dynamic ARP entries, and clears the dynamic ARP cache
IPSec	Configures IPSec
Neighbors	Displays IPv6 neighbor discovery protocol
Virtualization	Manages the virtualization and virtual machines
Virtual Switch Mgmt	Configures the system profile
Web	Configures web user interface and proxy settings
SNMP	Configures SNMP attributes, SNMP admin user, and trap sinks

Submenu Title	Description
Email Alerts	Configures the destination of email alerts and the recipients to be notified
XML gateway	Provides an XML request-response protocol to get and set hardware management information
JSON API	Manages JSON API
Logging	Sets up system log files, remote log sinks, and log formats
Configurations	Manages, activates, saves, and imports OS configuration files, and executes CLI commands
Docker	Manages docker images and containers.
Date and Time	Configures the date, time, and time zone of the switch system
NTP	Configures NTP (Network Time Protocol) and NTP servers
Licensing	Manages OS licenses

3.4.5 System Menu

The System menu makes available the following sub-menus (listed in order of appearance from top to bottom):

Submenu Title	Description
Modules	Displays a graphic illustration of the system modules. By moving the mouse over the ports in the front view, a pop-up caption is displayed to indicate the status of the port. The port state (active/down) is differentiated by a color scheme (green for active, gray/black for down). By moving the mouse over the rear view, a pop-up caption is displayed to indicate the leaf part information.
Inventory	Displays a table with the following information about the system modules: module name, type, serial number, ordering part number and ASIC firmware version
Power Management	Displays a table with the following information about the system power supplies: power supply name, power, voltage level, current consumption, and status. A total power summary table is also displayed providing the power used, the power capacity, and the power available.
OS Upgrade	Displays the installed OS images (and the active partition), uploads a new image, and installs a new image
Reboot	Reboots the system. Make sure that you save your configuration prior to clicking reboot.

3.4.6 Security Menu

The Security menu makes available the following submenus (listed in order of appearance from top to bottom):

Submenu Title	Description
Users	Manages (setting up, removing, modifying) user accounts
Admin Password	Modifies the system administrator password
SSH	Displays and generate host keys
AAA	Configures AAA (Authentication, Authorization, and Accounting) security services such as authentication methods and authorization
Login Attempts	Manages login attempts
RADIUS	Manages Radius client
TACACS+	Manages TACACS+ client
LDAP	Manages LDAP client
Certificate	Manages certificates

3.4.7 Ports Menu

The Ports menu displays the port state and enables some configuration attributes of a selected port. It also enables modification of the port configuration. A graphical display of traffic over time (last hour or last day) through the port is also available.

Submenu Title	Description
Ports	Manages port attributes, counters, transceiver info and displays a graphical counters histogram
Phy Profile	Provides the ability to manage PHY profiles
Monitor Session	Displays monitor session summary and enables configuration of a selected session
Telemetry	Displays and configures telemetry

3.4.8 Status Menu

The Status menu makes available the following submenus (listed in order of appearance from top to bottom):

Submenu Title	Description
Summary	Displays general information about the switch system and the OS image, including current date and time, hostname, uptime of system, system memory, CPU load averages, etc.
Profile and Capabilities	Displays general information about the switch system capabilities such as the enabled profiles (e.g IB/ETH) and their corresponding values
What Just Happened	Displays and configures What Just Happened packet drop reasons
Temperature	Provides a graphical display of the switch module sensors' temperature levels over time (1 hour). It is possible to display either the temperature level of one module's sensor or the temperature levels of all the module sensors' together.

Submenu Title	Description
Power Supplies	Provides a graphical display of one of the switch's power supplies voltage level over time (1 hour)
Fans	Provides a graphical display of fan speeds over time (1 hour). The display is per fan unit within a fan module.
CPU Load	Provides a graphical display of the management CPU load over time (1 hour)
Memory	Provides a graphical display of memory utilization over time (1 day)
Network	Provides a graphical display of network usage (transmitted and received packets) over time (1 day). It also provides per interface statistics.
Logs	Displays the system log messages. It is possible to display either the currently saved system log or a continuous system log.
Maintenance	Performs specific maintenance operations automatically on a predefined schedule
Alerts	Displays a list of the recent health alerts and enables the user to configure health settings
Virtualization	Displays the virtual machines, networks and volumes

3.4.9 ETH Mgmt Menu

The ETH Mgmt menu makes available the following sub-menus (listed in order of appearance from top to bottom):

Submenu Title	Description
Spanning Tree	Configures and monitors spanning tree protocol
MAC Table	Configures static mac addresses in the switch, and displays the MAC address table
Link Aggregation	Configures and monitors aggregated Ethernet links (LAG) and configures LACP
VLAN	Manages the switch VLAN table
MLAG	Manages multi-chassis LAGs
IGMP Snooping	Manages IGMP snooping in the switch
ACL	Manages Access Control in the switch
Priority Flow Control	Manages priority flow control
BGP	Manages the Border Gateway Protocol (BGP)
BGP IPv4	Displays the Border Gateway Protocol (BGP) IPv4 information
BGP IPv6	Displays the Border Gateway Protocol (BGP) IPv6 information

3.4.10 IP Route Menu

The IP Route menu makes available the following sub-menus (listed in order of appearance from top to bottom):

Submenu Title	Description
Router Global	Enables/disables IP routing protocol
IP Route	Configures, removes, and displays the routing table for router interfaces
IP Interface	Displays router interfaces
Address Resolution	Displays the address resolution (ARP) table for router interfaces
IP Diagnostic	Not implemented

3.5 UI Commands

3.5.1 CLI Session

- 3.5.1 CLI Session
  - 3.5.1.1 cli clear-history
  - 3.5.1.2 cli default
  - 3.5.1.3 cli max-sessions
  - 3.5.1.4 cli session
  - 3.5.1.5 terminal
  - 3.5.1.6 terminal sysrq enable
  - 3.5.1.7 show cli
  - 3.5.1.8 show cli max-sessions
  - 3.5.1.9 show cli num-sessions
  - 3.5.1.10 Banner
    - 3.5.1.10.1 banner login
    - 3.5.1.10.2 banner login-local
    - 3.5.1.10.3 banner login-remote
    - 3.5.1.10.4 banner logout
    - 3.5.1.10.5 banner logout-local
    - 3.5.1.10.6 banner logout-remote
    - 3.5.1.10.7 banner motd
    - 3.5.1.10.8 show banner
  - 3.5.1.11 SSH
    - 3.5.1.11.1 ssh server enable
    - 3.5.1.11.2 ssh server host-key
    - 3.5.1.11.3 ssh server listen
    - 3.5.1.11.4 ssh server login attempts
    - 3.5.1.11.5 ssh server login timeout

- 3.5.1.11.6 ssh server login record-period
- 3.5.1.11.7 ssh server min-version
- 3.5.1.11.8 ssh server ports
- 3.5.1.11.9 ssh server security strict
- 3.5.1.11.9 ssh server security strict
- 3.5.1.11.11 ssh server x11-forwarding
- 3.5.1.11.12 ssh client global
- 3.5.1.11.13 ssh client user
- 3.5.1.11.14 slogin
- 3.5.1.11.15 show ssh client
- 3.5.1.11.16 show ssh server
- 3.5.1.11.17 show ssh server host-keys
- 3.5.1.11.18 show ssh server login record-period

- 3.5.1.12 Remote Login
  - 3.5.1.12.1 telnet
  - 3.5.1.12.2 telnet-server enable
  - 3.5.1.12.3 show telnet-server

- 3.5.2 Web Interface
  - 3.5.2.1 web auto-logout
  - 3.5.2.2 web cache-enable
  - 3.5.2.3 web client cert-verify
  - 3.5.2.4 web client ca-list
  - 3.5.2.5 web enable
  - 3.5.2.6 web http
  - 3.5.2.7 web httpd
  - 3.5.2.8 web https
  - 3.5.2.9 web https ssl renegotiation enable
  - 3.5.2.10 web https ssl secure-cookie enable
  - 3.5.2.11 web proxy auth authtype
  - 3.5.2.12 web proxy auth basic
  - 3.5.2.13 web session timeout
  - 3.5.2.14 web session renewal
  - 3.5.2.15 show web

This section displays all the relevant commands used to manage CLI session terminal.

3.5.1.1 cli clear-history

	cli clear-history Clears the command history of the current user.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Example	`switch (config) # cli clear-history`

Related Commands	show cli
Notes	

3.5.1.2 cli default

	cli default {auto-logout <minutes> | paging enable | prefix-modes {enable | show-config} | progress enable | prompt {confirm-reload | confirm-reset | confirm-unsaved | empty-password}} no cli default {auto-logout | paging enable | prefix-modes {enable | show-config} | progress enable prompt {confirm-reload | confirm-reset | confirm-unsaved | empty-password} Configures default CLI options for this session only. The no form of the command deletes or disables the default CLI options.	
Syntax Description	auto-logout	Configures keyboard inactivity timeout for automatic logout. Range is 0-35791 minutes. Setting the value to 0 or using the no form of the command disables the auto-logout.
	paging enable	Enables text viewing one screen at a time.
	prefix-modes {enable | show-config}	Configures the prefix modes feature of CLI. • “prefix-modes enable” enables prefix modes for current session • “prefix-modes show-config” uses prefix modes in “show configuration” output for current session
	progress enable	Enables progress updates.
	prompt confirm-reload	Prompts for confirmation before rebooting.
	prompt confirm-reset	Prompts for confirmation before resetting to factory state.
	prompt confirm-unsaved	Confirms whether or not to save unsaved changes before rebooting.
	prompt empty-password	Prompts for a password if none is specified in a pseudo-URL for SCP.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # cli default prefix-modes enable`	
Related Commands	show cli	
Notes		

3.5.1.3 cli max-sessions

	cli max-sessions <number> no cli max-sessions Configures the maximum number of simultaneous CLI sessions allowed. The no form of the command resets this value to its default.	
Syntax Description	number	Range: 3-30

Default	30 sessions
Configuration Mode	config
History	3.5.0200
Example	`switch (config) # cli max-sessions 40`
Related Commands	show terminal
Notes	

3.5.1.4 cli session

	cli session {auto-logout <minutes> | paging enable | prefix-modes enable | progress enable | terminal {length <size> | resize | type <terminal-type> | width} | x-display full <display>} no cli session {auto-logout | paging enable | prefix-modes enable | progress enable | terminal type | x-display} Configures CLI options for this session only. The no form of the command deletes or disables the CLI sessions.	
Syntax Description	minutes	Configures keyboard inactivity timeout for automatic logout. Range: 0-35791 minutes Setting the value to 0 or using the no form of the command disables the auto logout.
	paging enable	Enables text viewing one screen at a time.
	prefix-modes enable	Configures the prefix modes feature of CLI and enables prefix modes for current session.
	progress enable	Enables progress updates.
	terminal length	Sets the number of lines for the current terminal. Range: 5-999
	terminal resize	Resizes the CLI terminal settings (to match the actual terminal window).
	terminal-type	Sets terminal type. Valid options are: • ansi • console • dumb • linux • unknown • vt52 • vt100 • vt102 • vt220 • xterm
	terminal width	Sets the width of the terminal in characters. Range: 34-999
	x-display full <display>	Specifies the display as a raw string (e.g. localhost:0.0).
Default	N/A	
Configuration Mode	config	
History	3.1.0000 3.8.2100: Removed "prefix-modes show-config" option and terminal type vt320	

Example	`switch (config) # cli session auto-logout`
Related Commands	show terminal
Notes	The "minutes" attribute can be configured from the CLI shell only.

3.5.1.5 terminal

	terminal {length <number of lines> | resize | type <terminal type> | width <number of characters>} no terminal type Configures default CLI options for this session only. The no form of the command clears the terminal type.	
Syntax Description	length	Sets the number of lines for this terminal. Range: 5-999
	resize	Resizes the CLI terminal settings (to match with real terminal).
	type	Sets the terminal type. Possible values: ansi, console, dumb, linux, screen, vt52, vt100, vt102, vt220, xterm.
	width	Sets the width of this terminal in characters. Range: 34-999
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # terminal length 500`	
Related Commands	show terminal	
Notes		

3.5.1.6 terminal sysrq enable

	terminal sysrq enable no terminal sysrq enable Enable SysRq over the serial connection (RS232 or Console port). The no form of the command disables SysRq over the serial connection (RS232 or Console port).
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.4.3000 3.9.3100: Updated command to be disabled by default
Example	`switch (config) # terminal sysrq enable`
Related Commands	show terminal
Notes	

3.5.1.7 show cli

	show cli Displays the CLI configuration and status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000

Example	`switch (config) # show cli`
	`CLI current session settings:`
	`Maximum line size: 8192`
	`Terminal width: 171 columns`
	`Terminal length: 38 rows`
	`Terminal type: xterm`
	`X display setting: (none)`
	`Auto-logout: disabled`
	`Paging: enabled`
	`Progress tracking: enabled`
	`Prefix modes: disabled`
	
	`CLI defaults for current session:`
	`Auto-logout: disabled`
	`Paging: enabled`
	`Progress tracking: enabled`
	`Prefix modes: enabled (and use in 'show configuration')`
	
	`Settings for current session:`
	`Show hidden config: yes`
	`Confirm losing changes: yes`
	`Confirm reboot/shutdown: no`
	`Confirm factory reset: yes`
	`Prompt on empty password: yes`

Related Commands	cli default
Notes	

3.5.1.8 show cli max-sessions

	show cli max-sessions Displays maximum number of sessions.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.5.0200
Example	`switch (config) # show cli max-sessions` `Maximum number of CLI sessions: 5`
Related Commands	
Notes	

3.5.1.9 show cli num-sessions

	show cli num-sessions Displays current number of sessions.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.5.0200
Example	`switch (config) # show cli num-sessions` `Current number of CLI sessions: 40`
Related Commands	
Notes	

3.5.1.10 Banner

3.5.1.10.1 banner login

	banner login <string> no banner login Sets the CLI welcome banner message. The no form of the command resets the system login banner to its default.
Syntax Description	N/A
Default	NVIDIA Onyx Switch Management
Configuration Mode	Any command mode
History	3.5.0200
Example	`switch (config) # banner login Example`
Related Commands	show banner
Notes	If more than one word is used (there is a space) quotation marks should be added (i.e., “xxxx xxxx”).

3.5.1.10.2 banner login-local

	banner login-local <string> no banner login-local Sets system login local banner. The no form of the command resets the banner to its default value.
Syntax Description	N/A
Default	""
Configuration Mode	Any command mode
History	3.1.0000 3.5.0200: Added the no form of the command
Example	`switch (config) # banner login-local Example`

Related Commands	show banner
Notes	• The login-local refers to the serial connection banner • If more than one word is used (there is a space) quotation marks should be added (i.e., “xxxx xxxx”)

3.5.1.10.3 banner login-remote

	banner login-remote <string> no banner login-remote Sets system login remote banner. The no form of the command resets the banner to its default value.	
Syntax Description	string	Text string
Default	""	
Configuration Mode	config	
History	3.1.0000 3.5.0200: Added the no form of the command	
Example	`switch (config) # banner login-remote Example`	
Related Commands	show banner	
Notes	• The login-remote refers to the SSH connections banner • If more than one word is used (there is a space) quotation marks should be added (i.e., “xxxx xxxx”).	

3.5.1.10.4 banner logout

	banner logout <string> no banner logout Sets system logout banner (for both local and remote logins). The no form of the command resets the banner to its default value.	
Syntax Description	string	Text string
Default	""	
Configuration Mode	config	
History	3.1.0000 3.5.0200: Added the no form of the command	
Example	`switch (config) # banner logout Example`	
Related Commands	show banner	
Notes	If more than one word is used (there is a space) quotation marks should be added (i.e., “xxxx xxxx”).	

3.5.1.10.5 banner logout-local

	banner logout-local <string> no banner logout-local Sets system logout local banner. The no form of the command resets the banner to its default value.

Syntax Description	string	Text string
Default	""	
Configuration Mode	config	
History	3.5.0200	
Example	`switch (config) # banner logout-local Example`	
Related Commands	show banner	
Notes	• The logout-local refers to the serial connection banner • If more than one word is used (there is a space) quotation marks should be added (i.e., “xxxx xxxx”).	

3.5.1.10.6 banner logout-remote

	banner logout-remote <string> no banner logout-remote Sets system logout remote banner. The no form of the command resets the banner to its default value.	
Syntax Description	string	Text string
Default	""	
Configuration Mode	config	
History	3.5.0200	
Example	`switch (config) # banner logout-remote Example`	
Related Commands	show banner	
Notes	• The logout-remote refers to SSH connections banner • If more than one word is used (there is a space) quotation marks should be added (i.e., “xxxx xxxx”).	

3.5.1.10.7 banner motd

	banner motd <string> no banner motd Configures the message of the day banner. The no form of the command resets the system Message of the Day banner.	
Syntax Description	string	Text string
Default	NVIDIA Switch	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # banner motd “My Banner”`	
Related Commands	show banner	
Notes	• If more than one word is used (there is a space) quotation marks should be added (i.e., “xxxx xxxx”). • To insert a multi-line MotD, hit Ctrl-V (escape sequence) followed by Ctrl-J (new line sequence). The symbol “^J” should appear. Then, whatever is typed after it becomes the new line of the MotD. Remember to also include the string between quotation marks.	

3.5.1.10.8 show banner

	show banner Sets system logout remote banner. The no form of the command resets the banner to its default value.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.5.0200	Updated example
	3.6.6000	Updated example
	3.9.3200	Updated example

Example	`switch (config) # show banner Banners: Message of the Day (MOTD):`
	`Switch Login:`
	`NVIDIA ONYX Switch Management`
	`Logout: Goodbye`	

Related Commands	banner login banner login-local banner login-remote banner logout banner logout-local banner logout-remote banner motd	
Notes		

3.5.1.11 SSH

3.5.1.11.1 ssh server enable

	ssh server enable no ssh server enable Enables the SSH server. The no form of the command disables the SSH server.
Syntax Description	N/A
Default	SSH server is enabled
Configuration Mode	config
History	3.1.0000
Example	`switch (config) # ssh server enable`
Related Commands	show banner
Notes	Disabling SSH server does not terminate existing SSH sessions, it only prevents new ones from being established.

3.5.1.11.2 ssh server host-key

	ssh server host-key {<key-type> {private-key <private-key>| public-key <public-key>} | generate} Configures host keys for SSH.	
Syntax Description	key-type	• rsa2—RSAv2 • dsa2—DSAv2
	private-key	Sets new private-key for the host keys of the specified type.
	public-key	Sets new public-key for the host keys of the specified type.
	generate	Generates new RSA and DSA host keys for SSH.
Default	SSH keys are locally generated	
Configuration Mode	config	
History	3.1.0000 3.4.2300: Added notes 3.9.0300: Removed RSAv1 3.9.1000: Added a note	

Example	`switch (config) # ssh server host-key dsa2 private-key`
	`Key: **********`
	`Confirm: **********`	

Related Commands	show banner	
Notes	RSA2 and a DSA2 host keys are generated by default. The RSA2 key can be used as SSH server and client, while DSA2 key can only be used as SSH client. When the switch is a server, use RSA key to connect to the NVIDIA Onyx device. When the switch is a client (e.g. downloading image or uploading logs), RSA key is recommended. DSA key is only for legacy devices and has been deprecated by OpenSSH starting with the 7.0 release.	

3.5.1.11.3 ssh server listen

	ssh server listen {enable | interface <inf>} no ssh server listen {enable | interface <inf>} Enables the listen interface restricted list for SSH. If enabled, and at least one non-DHCP interface is specified in the list, the SSH connections are only accepted on those specified interfaces. The no form of the command disables the listen interface restricted list for SSH. When disabled, SSH connections are not accepted on any interface.	
Syntax Description	enable	Enables SSH interface restrictions on access to this system.
	interface	Adds interface to SSH server access restriction list. Possible interfaces are “lo”, and “mgmt0”.
Default	SSH listen is enabled	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # ssh server listen enable`	
Related Commands	show ssh server	
Notes		

3.5.1.11.4 ssh server login attempts

	ssh server login attempts <number> no ssh server login attempts Configures maximum login attempts on SSH server. The no form of the command resets the login attempts value to its default.	
Syntax Description	number	Range: 3-100 attempts
	interface	Adds interface to SSH server access restriction list. Possible interfaces are “lo”, and “mgmt0”.
Default	6 attempts	
Configuration Mode	config	
History	3.1.0000 3.5.1000: Increased minimum number of attempts 3.9.0900: Added notes	
Example	`switch (config) # ssh server login attempts 5`	
Related Commands	show ssh server	
Notes	• The number configured with this command will be relevant only if it is equal or smaller than the number of password prompts • Be aware that the "aaa authentication attempts lockout max-fail" default is 5, and the user might be locked before this command will have an affect. Both numbers need to be configured	

3.5.1.11.5 ssh server login timeout

	ssh server login timeout <time> no ssh server login timeout Configures login timeout on SSH server. The no form of the command resets the timeout value to its default.	
Syntax Description	time	Range: 1-600 seconds
Default	120 seconds	
Configuration Mode	config	
History	3.5.0200	
Example	`switch (config) # ssh server login timeout 130`	
Related Commands	show ssh server	
Notes		

3.5.1.11.6 ssh server login record-period

	ssh server login record-period <days> no ssh server login record-period Configures the amount of days for counting the number of successful logins. The no form of the command disabled this function.	
Syntax Description	Days	Range: 1-30 days Default: 1 day
Default	Disabled	

Configuration Mode	config
History	3.9.0300 3.9.0500: Changed "SSH server login record-period" default value to 1 day
Example	`switch (config) # ssh server login record-period 1`
Related Commands	show ssh server login record-period show ssh server
Notes	

3.5.1.11.7 ssh server min-version

	ssh server min-version <version> no ssh server min-version Sets the minimum version of the SSH protocol that the server supports. The no form of the command resets the minimum version of SSH protocol supported.	
Syntax Description	version	Possible versions are 1 and 2
Default	2	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # ssh server min-version 2`	
Related Commands	show ssh server	
Notes		

3.5.1.11.8 ssh server ports

	ssh server ports {<port1> [<port2>...]} Specifies which ports the SSH server listens on.	
Syntax Description	port	Port number between [1-65535]
Default	22	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # ssh server ports 22`	
Related Commands	show ssh server	
Notes	• Multiple ports can be specified by repeating the <port> parameter • The command will remove any previous ports if not listed in the command	

3.5.1.11.9 ssh server security strict

	ssh server ports {<port1> [<port2>...]} Enables strict security settings. The no form of the command disables strict security settings.
Syntax Description	N/A
Default	N/A

Configuration Mode	config
History	3.3.5060 3.6.4000 3.9.0300: Updated notes
Example	`switch (config) # ssh server security strict`
Related Commands	show ssh server
Notes	The following ciphers are disabled for SSH when strict security is enabled: • aes256-cbc • aes192-cbc • aes128-cbc • rijndael-cbc@lysator.liu.se • 3des-cbc

3.5.1.11.10 ssh server security strict

	ssh server tcp-forwarding enable Enables TCP port forwarding. The no form of the command disables TCP port forwarding.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Example	`switch (config) # ssh server tcp-forwarding enable`
Related Commands	show ssh server
Notes	

3.5.1.11.11 ssh server x11-forwarding

	ssh server x11-forwarding enable no ssh server x11-forwarding enable Enables X11 forwarding on the SSH server. The no form of the command disables X11 forwarding.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.1.0000
Example	`switch (config) # ssh server x11-forwarding enable`
Related Commands	
Notes	

3.5.1.11.12 ssh client global

	ssh client global {host-key-check <policy>} | known-host <known-host-entry>} no ssh client global {host-key-check | known-host localhost} Configures global SSH client settings. The no form of the command negates global SSH client settings.	
Syntax Description	host-key-check <policy>	Sets SSH client configuration to control how host key checking is performed. This parameter may be set in 3 ways. • If set to “no” it always permits connection, and accepts any new or changed host keys without checking • If set to “ask” it prompts user to accept new host keys, but does not permit a connection if there was already a known host entry that does not match the one presented by the host • If set to “yes” it only permits connection if a matching host key is already in the known hosts file
	known-host	Adds an entry to the global known-hosts configuration file
	known-host-entry	Adds/removes an entry to/from the global known-hosts configuration file. The entry consist of “<IP> <key-type> <key>”.
Default	host-key-check - ask, no keys are configured by default	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # ssh client global host-key-check no` `switch (config) # ssh client global known-host "72.30.2.2 ssh-rsa AAAAB3NzaC1yc2EAAAAB....f2CyXFq4pzaR1jar1Vk="`	
Related Commands	show ssh client	
Notes		

3.5.1.11.13 ssh client user

	ssh client user <username> {authorized-key sshv2 <public key> | identity <key type> {generate | private-key [<private key>] | public-key [<public key>]} | known-host <known host> remove} no ssh client user admin {authorized-key sshv2 <public key ID> | identity <key type>} Adds an entry to the global known-hosts configuration file, either by generating new key, or by adding manually a public or private key. The no form of the command removes a public key from the specified user's authorized key list, or changes the key type.	
Syntax Description	username	The specified user must be a valid account on the system. Possible values for this parameter are “admin”, “monitor”, “xmladmin”, and “xmluser”.
	authorized-key sshv2 <public key>	Adds the specified key to the list of authorized SSHv2 RSA or DSA public keys for this user account. These keys can be used to log into the user's account.
	identity <key type>	Sets certain SSH client identity settings for a user, dsa2 or rsa2.
	generate	Generates SSH client identity keys for specified user.
	private-key	Sets private key SSH client identity settings for the user.
	public-key	Sets public key SSH client identity settings for the user.

	known-host <known host> remove	Removes host from user's known host file.
Default	No keys are created by default	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # ssh client user admin known-host 172.30.1.116 remove`	
Related Commands	show ssh client	
Notes	If a key is being pasted from a cut buffer and was displayed with a paging program, it is likely that newline characters have been inserted, even if the output was not long enough to require paging. One can specify “no cli session paging enable” before running the “show” command to prevent the newlines from being inserted.	

3.5.1.11.14 slogin

	slogin [<slogin options>] <hostname> Invokes the SSH client. The user is returned to the CLI when SSH finishes.		
Syntax Description	slogin options	-p -c -L -l -m -R -o -1 -2 -4 -6 -g -q -V -v -x -X -Y -y -a -A	-o flags (option allowed flags): AdressFamily BatchMode CheckHostIP Cipher Ciphers ConnectTimeout ForwardAgent ForwardX11 ForwardX11Trusted HostKeyAlgorithms KexAlgorithms LogLevel MACs Port PubkeyAcceptedKeyTypes PubkeyAuthentication StrictHostKeyChecking TCPKeepAlive User VerifyHostKeyDNS
	vrf_name	There are no restrictions on the VRF name, as long as the VRF exists in the switch.	
Default	N/A		
Configuration Mode	config		
History	3.1.0000 3.10.1000: Updated the slogin options		

Example	`switch (config) # slogin 192.168.10.70`
	`The authenticity of host '192.168.10.70 (192.168.10.70)' can't be established.`
	`RSA key fingerprint is 2e:ad:2d:23:45:4e:47:e0:2c:ae:8c:34:f0:1a:88:cb.`
	`Are you sure you want to continue connecting (yes/no)? yes`		

Related Commands			

Notes	For more information about slogin options see the following: linux.die.net/man/1/ssh

3.5.1.11.15 show ssh client

	show ssh client Displays the client configuration of the SSH server.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000

Example	`switch (config) # show ssh client`
	`SSH client Strict Hostkey Checking: ask`
	
	`SSH Global Known Hosts:`
	`Entry 1: 72.30.2.2`
	`Finger Print: 1e:b7:8b:ec:ab:35:98:be:6b:d6:12:c2:18:72:12:d6`
	
	`No SSH user identities configured.`
	
	`No SSH authorized keys configured.`

Related Commands	
Notes	

3.5.1.11.16 show ssh server

	show ssh server Displays SSH server configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.4.0000: Updated example 3.5.0200: Added SSH login timeout and max attempts 3.6.6000: Updated example 3.9.0300: Updated example—removed RSA v1 and added SSH server login record-period 3.9.0500: Changed "SSH server login record-period" default period to 1 day

Example	switch (config) # show ssh server SSH server configuration: SSH server enabled: yes Server security strict mode: no Minimum protocol version: 2 TCP forwarding enabled: yes X11 forwarding enabled: no SSH login timeout: 120 SSH login max attempts: 6 SSH server login record-period: 1 SSH server ports: 22 Interface listen enabled: yes Listen Interfaces: No interface configured. Host Key Finger Prints and Key Lengths: RSA v2 host key: SHA256:gVu6qLW1ZifEp8wRer2jkvILZMGNl6VCYU3HqC1INC8 (2048) DSA v2 host key: SHA256:JnldTEla20ZF/c5LdIqo9251DzO742k3hFCQh3Jt4ZA (1024)
Related Commands	
Notes	

3.5.1.11.17 show ssh server host-keys

	show ssh server host-keys Displays SSH host key configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.6.6000: Updated example 3.9.0300: Updated example—removed RSA v1
Example	switch (config) # show ssh server host-keys SSH server configuration: SSH server enabled: yes Server security strict mode: no Minimum protocol version: 2 TCP forwarding enabled: yes X11 forwarding enabled: no SSH login timeout: 120 SSH login max attempts: 6 SSH server ports: 22 Interface listen enabled: yes Listen Interfaces: No interface configured. Host Key Finger Prints and Key Lengths: RSA v2 host key: SHA256:gVu6qLW1ZifEp8wRer2jkvILZMGNl6VCYU3HqC1INC8 (2048) DSA v2 host key: SHA256:JnldTEla20ZF/c5LdIqo9251DzO742k3hFCQh3Jt4ZA (1024) Host Keys: RSA v2 host key: "kebo-2100-1 ssh-rsa AAAAB3Nza<...>KE5" DSA v2 host key: "kebo-2100-1 ssh-dss AAAAB3Nza<...>/s="
Related Commands	ssh server host-keys
Notes	

3.5.1.11.18 show ssh server login record-period

	show ssh server login record-period Displays the amount of days for counting the number of successful logins.(Default: 30 days)
Syntax Description	N/A
Default	Disabled
Configuration Mode	Any command mode
History	3.9.0300 3.9.0500: Changed "SSH server login record-period" default value to 1 day
Example	switch (config) # show ssh server login record-period SSH server login record-period: 1
Related Commands	ssh server login record-period
Notes	

3.5.1.12 Remote Login

3.5.1.12.1 telnet

	telnet Logs into another system using telnet.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Example	`switch (config) # telnet` `telnet>`
Related Commands	telnet-server
Notes	

3.5.1.12.2 telnet-server enable

	telnet-server enable no telnet-server enable Enables the telnet server. The no form of the command disables the telnet server.
Syntax Description	N/A
Default	Telnet server is disabled
Configuration Mode	config
History	3.1.0000
Example	`switch (config) # telnet-server enable`

Related Commands	telnet-server show telnet-server
Notes	

3.5.1.12.3 show telnet-server

	show telnet-server Displays telnet server settings.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Example	`switch (config) # show telnet-server` `Telnet server enabled: yes`
Related Commands	telnet-server show telnet-server
Notes	

3.5.2 Web Interface

3.5.2.1 web auto-logout

	web auto-logout <mins> no web auto-logout <mins> Configures length of user inactivity before auto-logout of a web session. The no form of the command disables the web auto-logout (web sessions will never logged out due to inactivity).	
Syntax Description	mins	The length of user inactivity in minutes "0" disables the inactivity timer (same as a “no web auto-logout” command)
Default	60 minutes	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # web auto-logout 60`	
Related Commands	show web	
Notes	The no form of the command does not automatically log users out due to inactivity.	

3.5.2.2 web cache-enable

	web cache-enable no web cache-enable Enables web clients to cache web pages. The no form of the command disables web clients from caching web pages.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config
History	3.4.1100
Example	`switch (config) # no web cache-enable`
Related Commands	show web
Notes	

3.5.2.3 web client cert-verify

	web client cert-verify no web client cert-verify Enables verification of server certificates during HTTPS file transfers. The no form of the command disables verification of server certificates during HTTPS file transfers.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.2.3000
Example	`switch (config) # web client cert-verify`
Related Commands	
Notes	

3.5.2.4 web client ca-list

	web client ca-list {<ca-list-name> | default-ca-list | none} no web client ca-list Configures supplemental CA certificates for verification of server certificates during HTTPS file transfers. The no form of the command uses no supplemental certificates.	
Syntax Description	ca-list-name	Specifies CA list to configure
	default-ca-list	Configures default supplemental CA certificate list
	none	Uses no supplemental certificates
Default	default-ca-list	
Configuration Mode	config	
History	3.2.3000	

Example	`switch (config) # web client ca-list default-ca-list`
Related Commands	
Notes	

3.5.2.5 web enable

	web [vrf <vrf-name>] enable [force] no web [vrf <vrf-name>] enable Enables the web-based management console. The no form of the command disables the web-based management console.
Syntax Description	vrf name—Describes VRF name for web daemon. If the VRF parameter is not specified, the "default" VRF will be used implicitly force—Restarts web with passed VRF context even if it was already enabled using other VRF
Default	enable
Configuration Mode	config
History	3.1.0000 3.8.1000—Added note 3.9.2000—Added VRF option
Example	`switch (config) # web enable`
Related Commands	show web
Notes	Web interface can be enabled only in one VRF at a time.

3.5.2.6 web http

	web http {enable | port <port-number> | redirect} no web http {enable | port | redirect} Configures HTTP access to the web-based management console. The no form of the command negates HTTP settings for the web-based management console.	
Syntax Description	enable	Enables HTTP access to the web-based management console.
	port-number	Sets a port for HTTP access.
	redirect	Enables redirection to HTTPS. If HTTP access is enabled, this specifies whether a redirect from the HTTP port to the HTTPS port should be issued to mandate secure HTTPS access.
Default	• HTTP is disabled • HTTP TCP port is 80 • HTTP redirect to HTTPS is disabled	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # web http enable`	
Related Commands	show web web enable	

Notes	Enabling HTTP is meaningful if the WebUI as a whole is enabled

3.5.2.7 web httpd

	web httpd listen {enable | interface <ifName>} no web httpd listen {enable | interface <ifName>} Enables the listen interface restricted list for HTTP and HTTPS. The no form of the command disables the HTTP server listen ability.	
Syntax Description	enable	Enables Web interface restrictions on access to this system.
	interface <ifName>	Adds interface to Web server access restriction list (i.e., mgmt0, mgmt1).
Default	• Listening is enabled • All interfaces are permitted.	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # web httpd listen enable`	
Related Commands	show web web enable	
Notes	If enabled, and if at least one of the interfaces listed is eligible to be a listen interface, then HTTP/HTTPS requests will only be accepted on those interfaces. Otherwise, HTTP/HTTPS requests are accepted on any interface.	

3.5.2.8 web https

	web https {certificate {regenerate | name | default-cert} | enable | port <port number> | ssl ciphers {all | TLS | TLS1.2}} no web https {enable | port <port number>} Configures HTTPS access to the web-based management console. The no form of the command negates HTTPS settings for the web-based management console.	
Syntax Description	certificate regenerate	Re-generates certificate to use for HTTPS connections
	certificate name	Configure the named certificate to be used for HTTPS connections
	certificate default-cert	Configure HTTPS to use the configured default certificate
	enable	Enables HTTPS access to the web-based management console
	port	Sets a TCP port for HTTPS access
	ssl ciphers {all | TLS | TLS1.2}	Sets ciphers to be used for HTTPS
Default	• HTTPS is enabled • Default port is 443	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Added “ssl ciphers” parameter

	3.4.0010	Added TLS parameter to “ssl ciphers”
	3.8.1000	Added note
Example	`switch (config) # web https enable`	
Related Commands	show web web enable	
Notes	• Enabling HTTPS is meaningful if the WebUI as a whole is enabled • See the command “crypto certificate default-cert name” for how to change the default certificate if inheriting the configured default certificate is preferred	

3.5.2.9 web https ssl renegotiation enable

	web https ssl renegotiation enable no web https ssl renegotiation enable Enables SSL renegotiation flag in httpd web server. The no form of the command disables SSL renegotiation flag in httpd web server.
Syntax Description	N/A
Default	• HTTPS is enabled • Default port is 443
Configuration Mode	config
History	3.6.8008
Example	`switch (config) # web https ssl renegotiation enable`
Related Commands	show web web enable
Notes	

3.5.2.10 web https ssl secure-cookie enable

	web https ssl secure-cookie enable no web https ssl secure-cookie enable Enables SSL secure-cookie flag in httpd web server. The no form of the command disables secure-cookie flag in httpd web server.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config
History	3.6.8008
Example	`switch (config) # web https ssl secure-cookie enable`
Related Commands	show web web enable
Notes	

3.5.2.11 web proxy auth authtype

	web proxy auth authtype <auth-type> no web proxy auth authtype Configures type of authentication to use with web proxy. The no form of the command resets web proxy authentication type to its default.	
Syntax Description	auth-type	Possible values: • none - no authentication • basic - HTTP basic authentication
Default	Basic authentication settings	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # web proxy auth authtype basic`	
Related Commands	show web web enable	
Notes		

3.5.2.12 web proxy auth basic

	web proxy auth basic {password <password> | username <username>} no web proxy auth basic {password | username} Configures HTTP basic authentication settings for proxy. The no form of the command clears password or username configuration.	
Syntax Description	password	Sets plaintext password for HTTP basic authentication with web proxy
	username	Sets username for HTTP basic authentication with web proxy
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # web proxy auth basic password 57R0ngP455w0rD`	
Related Commands	show web web enable	
Notes		

3.5.2.13 web session timeout

	web session timeout <number of minutes> Configures time after which a session expires	
Syntax Description	number of minutes	Number of minutes
Default	2 hr 30 min	
Configuration Mode	config	
History	3.1.0000	

Example	`switch (config) # web session timeout 180`
Related Commands	
Notes	

3.5.2.14 web session renewal

	web session renewal <number of minutes> Configures time before expiration to renew a session	
Syntax Description	number of minutes	Number of minutes
Default	30 min	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # web session renewal 20`	
Related Commands		
Notes		

3.5.2.15 show web

	show web Displays WebUI configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.6000 3.6.8008—Updated example 3.9.2000—Updated example, adding VRP field

Example	switch (config) # show web Web User Interface: Web interface enabled: yes VRF name: mgmt Web caching enabled: no HTTP enabled: no HTTP port: 80 HTTP redirect to HTTPS: no HTTPS enabled: yes HTTPS port: 443 HTTPS ssl-ciphers: TLS1.2 HTTPS ssl-renegotiation: no HTTPS ssl-secure-cookie: yes HTTPS certificate name: default-cert Listen enabled: yes Listen Interfaces: No interface configured. Inactivity timeout: 1 hr Session timeout: 2 hr 30 min Session renewal: 30 min Web file transfer proxy: Proxy enabled: no Web file transfer certificate authority: HTTPS server cert verify: yes HTTPS supplemental CA list: default-ca-list
Related Commands	web auto-logout web cache-enable web enable web http web httpd web https web https ssl renegotiation enable web https ssl secure-cookie enable web proxy auth authtype web proxy auth basic
Notes	

4 System Management

The following pages provide information on configuring general management features on the system.

- Management Interfaces
- Chassis Management
- Management Source IP Address
- Upgrade/Downgrade Process
- Configuration Management
- Resource Scale

4.1 Management Interfaces

Management interfaces are used in order to provide access to management user interfaces. NVIDIA switches support out-of-band (OOB) dedicated interfaces (e.g. mgmt0, mgmt1) and in-band dedicated interfaces. In addition, most systems feature a serial port that provides access to the CLI only.On systems with two OOB management ports, both of them may be configured on the same VLAN if needed. In this case, ARP replies to the IP of those management interfaces is answered from either of them.

4.1.1 Configuring Management Interfaces with Static IP Addresses

If the system was set during initialization to obtain dynamic IP addresses through DHCP and you wish to switch to static assignments, perform the following steps:

1. Enter Config configuration mode. Run:

```
switch > enable
switch # configure terminal
```

2. Disable setting IP addresses using the DHCP using the following command:

```
switch (config) # no interface <ifname> dhcp
```

3. Define your interfaces statically using the following command:

```
switch (config) # interface <ifname> ip address <IP address> <netmask>
```

4.1.2 Configuring IPv6 Address on the Management Interface

1. Enable IPv6 on this interface. Run:

```
switch (config) # interface mgmt0 ipv6 enable
```

2. Set the IPv6 address to be configured automatically. Run:

```
switch (config) # interface mgmt0 ipv6 address autoconfig
```

3. Verify the IPv6 address is configured correctly. Run:

```
switch (config) # show interfaces mgmt0 brief
```

4.1.3 Dynamic Host Configuration Protocol (DHCP)

DHCP is used for automatic retrieval of management IP addresses.

For all other systems (and software versions) DHCP is disabled by default.

If a user connects through SSH, runs the wizard and turns off DHCP, the connection is immediately terminated as the management interface loses its IP address.

```
<localhost># ssh admin@<ip-address>
NVIDIA Onyx Switch Management
Password:
NVIDIA switch
NVIDIA configuration wizard
Do you want to use the wizard for initial configuration? yes
Step 1: Hostname? [my-switch]
Step 2: Use DHCP on mgmt0 interface? [yes] no
<localhost>#
```

In this case the serial connection should be used.

4.1.4 Default Gateway

To configure manually the default gateway, use the “ip route” command, with “0.0.0.0” as prefix and mask. The next-hop address must be within the range of one of the IP interfaces on the system.

```
switch (config)# ip route 0.0.0.0 0.0.0.0 10.10.0.2
switch (config)# show ip route
Destination       Mask             Gateway         Interface     Source       Distance/Metric
default           0.0.0.0          10.10.0.2       mgmt0         static       0/0
10.10.0.0         255.255.254.0    0.0.0.0         mgmt0         direct       0/0
```

4.1.5 In-Band Management

In-band management is a management path passing through the data ports. In-band management can be created over one of the VLANs in the systems.

The in-band management feature does not require any license. However, it works only for the system profile Ethernet. It can be enabled with IP Routing.

To set an in-band management channel:

1. Create a VLAN. Run:

```
switch (config)# vlan 10
switch (config vlan 10)#
```

2. Create a VLAN interface. Run:

```
switch (config)# interface vlan 10
switch (config interface vlan 10)#
```

3. Configure L3 attributes on the newly created VLAN interface. Run:

```
switch (config interface vlan 10)# ip address 10.10.10.10 /24
```

4. (Optional) Verify in-band management configuration. Run:

```
switch (config)# show interfaces vlan 10
  Admin state: Enabled
  Operational state: Up
  Mac Address: f4:52:14:67:07:e8
  Internet Address: 10.10.10.10/24
  Broadcast address: 10.10.10.255
  MTU: 1500 bytes
  Arp timeout: 1500 seconds
  Icmp redirect: Disabled
  Description: N/A
  VRF: default
  Counters: Enabled
 RX
  0 Unicast packets
  0 Multicast packets
  0 Unicast bytes
  0 Multicast bytes
  0 Bad packets
  0 Bad bytes
 TX
  0 Unicast packets
  0 Multicast packets
  0 Unicast bytes
  0 Multicast bytes
```

4.1.6 Configuring Hostname via DHCP (DHCP Client Option 12)

This feature, also known as the DHCP Client Option 12, is enabled by default and assigns the switch system a hostname via DHCP as long as network manager configures hostname to the management interfaces' (i.e. mgmt0, mgmt1) MAC address. If a network manager configures the hostname manually through any of the user interfaces, the hostname is not retrieved from the DHCP server.

To enable fetching hostname from DHCP server, run:

```
switch (config interface mgmt0) # dhcp hostname
```

To disable fetching hostname from DHCP server, run:

```
switch (config interface mgmt0) # no dhcp hostname
```

> Getting the hostname through DHCP is enable by default and will change the switch hostname if the hostname is not set by the user. Therefore, if a switch is part of an HA cluster the user would need to make sure the HA master has the same HA node names as the DHCP server.

4.1.7 Management VRF

Management VRF is a virtual routing function that is responsible for providing IP services for switch management. It is the only VRF where outband management interface mgmt0 belongs.

Initially, a system has only one VRF—the default VRF. This VRF supports both management and data forwarding functions. A management VRF can them be created—mgmt and user VRFs (mgmt VRF is not created with image upgrade automatically). The mgmt VRF is also created on reset factory flows.

When mgmt VRF is created, all mgmt interfaces are automatically moved to it. New management functions can be shutdown in a default VRF and created in the management VRF. Also, management services can be started in 'user' VRFs, with the only difference that the 'user' VRF does not have mgmt interfaces.

Following services are considered management services:

Service	Run by VRF Once mgmt VRF is Created
ssh	Single instance in all VRFs
snmp-agent	Single instance in any VRF
snmptrap	Can be configured in multiple VRFs at the same time
syslogd	Can be configured in multiple VRFs at the same time
web server	Single instance in any VRF
ntp	Single instance in any VRF
dns	Single instance in any VRF
tacacs radius	Single instance in any VRF
OpenFlow API	Mgmt/default if mgmt is not created
sFlow	Single instance in any VRF
ftp-server	Mgmt/default if mgmt is not created
telnet-server	Mgmt/default if mgmt is not created
docker	Single instance in any VRF
ip filters	Single instance in all VRFs
ZTP	Mgmt only
IPL	Default VRF only

User VRF will have routing functions and its primary purpose is to perform routing of user traffic.

Default VRF is used to run some non-management system functions and can also be used to serve as a global routing instance for multi-VRF traffic.

When a service is moved from VRF to VRF, its configuration is removed.

4.1.8 Management Interface Commands

- 4.1.8.1 Interface
  - 4.1.8.1.1 interface
  - 4.1.8.1.2 ip address
  - 4.1.8.1.3 ip default-gateway
  - 4.1.8.1.4 alias
  - 4.1.8.1.5 mtu
  - 4.1.8.1.6 duplex
  - 4.1.8.1.7 speed
  - 4.1.8.1.8 dhcp
  - 4.1.8.1.9 dhcp hostname
  - 4.1.8.1.10 shutdown
  - 4.1.8.1.11 zeroconf
  - 4.1.8.1.12 comment
  - 4.1.8.1.13 ipv6 enable
  - 4.1.8.1.14 ipv6 address
  - 4.1.8.1.15 ipv6 dhcp primary-intf
  - 4.1.8.1.16 ipv6 dhcp stateless
  - 4.1.8.1.17 ipv6 dhcp client enable
  - 4.1.8.1.18 ipv6 dhcp client renew
  - 4.1.8.1.19 show interfaces mgmt0
  - 4.1.8.1.20 show interfaces mgmt0 brief
  - 4.1.8.1.21 show interfaces mgmt0 configured
- 4.1.8.2 Hostname Resolution
  - 4.1.8.2.1 hostname
  - 4.1.8.2.2 ip name-server
  - 4.1.8.2.3 ip domain-list
  - 4.1.8.2.4 ip/ipv6 host
  - 4.1.8.2.5 ip/ipv6 map-hostname
  - 4.1.8.2.6 show hosts
- 4.1.8.3 Routing
  - 4.1.8.3.1 IP route
  - 4.1.8.3.2 ipv6 default-gateway
  - 4.1.8.3.3 show ip/ipv6 route
  - 4.1.8.3.4 show ipv6 default-gateway
- 4.1.8.4 Network to Media Resolution (ARP & NDP)
  - 4.1.8.4.1 ip arp
  - 4.1.8.4.2 ip arp timeout
  - 4.1.8.4.3 ip arp cache-update
  - 4.1.8.4.4 show ip arp
  - 4.1.8.4.5 ipv6 neighbor
  - 4.1.8.4.6 clear ipv6 neighbors
  - 4.1.8.4.7 show ipv6 neighbors
- 4.1.8.5 DHCP
  - 4.1.8.5.1 ip dhcp
  - 4.1.8.5.2 show ip dhcp
- 4.1.8.6 General IPv6
  - 4.1.8.6.1 ipv6 enable

- 4.1.8.7 IP Diagnostic Tools
  - 4.1.8.7.1 ping
  - 4.1.8.7.2 traceroute
  - 4.1.8.7.3 tcpdump

4.1.8.1 Interface

4.1.8.1.1 interface

	interface {mgmt0 | mgmt1 | lo | vlan<id>} Enters a management interface context.	
Syntax Description	mgmt0	Management port 0 (out of band).
	mgmt1	Management port 1 (out of band).
	lo	Loopback interface
	vlan<id>	In-band management interface (e.g., vlan10)
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config)# interface mgmt0` `switch (config interface mgmt0)#`	
Related Commands	show interfaces <ifname>	
Notes		

4.1.8.1.2 ip address

	ip address <IP address> <netmask> no ip address Sets the IP address and netmask of this interface. The no form of the command clears the IP address and netmask of this interface.	
Syntax Description	IP address	IPv4 address
	netmask	Subnet mask of IP address
Default	0.0.0.0/0	
Configuration Mode	config interface management	
History	3.1.0000	
Example	`switch (config interface mgmt0)# ip address 10.10.10.10 255.255.255.0`	
Related Commands	show interfaces <ifname>	
Notes	If DHCP is enabled on the specified interface, then the DHCP IP assignment will hold until DHCP is disabled	

4.1.8.1.3 ip default-gateway

	ip default-gateway <next-hop-IP-address> <interface-name> no default-gateway <next-hop-IP-address> <interface-name> Configures a default route. The no form of the command removes the current default route.	
Syntax Description	next hop IP address	gateway IP address
	interface name	default gateway interface name
Default	N/A	
Configuration Mode	config interface management	
History	3.1.0000 3.8.1000: Updated Command & Syntax description	
Example	`switch (config interface mgmt0)# ip default-gateway mgmt1`	
Related Commands		
Notes		

4.1.8.1.4 alias

	alias <index> ip address < IP address> <netmask> no alias <index> Adds an additional IP address to the specified interface. The secondary address will appear in the output of “show interface” under the data of the primary interface along with the alias. The no form of the command removes the secondary address to the specified interface.	
Syntax Description	index	A number that is to be aliased to (associated with) the secondary IP.
	IP address	Additional IP address.
	netmask	Subnet mask of the IP address.
Default	N/A	
Configuration Mode	config interface management	
History	3.1.0000	
Example	`switch (config interface mgmt0)# alias 2 ip address 9.9.9.9 255.255.255.255`	
Related Commands	show interfaces <ifname>	
Notes	• If DHCP is enabled on the specified interface, then the DHCP IP assignment will hold until DHCP is disabled • More than one additional IP address can be added to the interface	

4.1.8.1.5 mtu

	mtu <bytes> no mtu <bytes> Sets the Maximum Transmission Unit (MTU) of this interface. The no form of the command resets the MTU to its default.	
Syntax Description	bytes	The entry range is 68-1500.
Default	1500	
Configuration Mode	config interface management	
History	3.6.3004	
Example	`switch (config interface mgmt0)# mtu 1500`	
Related Commands	show interfaces <ifname>	
Notes		

4.1.8.1.6 duplex

	duplex <duplex> no duplex Sets the interface duplex. The no form of the command resets the duplex setting for this interface to its default value.	
Syntax Description	duplex	Sets the duplex mode of the interface. The following are the possible values: • half—half duplex • full—full duplex • auto—auto duplex sensing (half or full)
Default	auto	
Configuration Mode	config interface management	
History	3.1.0000	
Example	`switch (config interface mgmt0)# duplex auto`	
Related Commands	show interfaces <ifname>	
Notes	• Setting the duplex to “auto” also sets the speed to “auto” • Setting the duplex to one of the settings “half” or “full” also sets the speed to a manual setting which is determined by querying the interface to find out its current auto-detected state	

4.1.8.1.7 speed

	speed <speed> no speed Sets the interface speed. The no form of the command resets the speed setting for this interface to its default value.

Syntax Description	speed	Sets the speed of the interface. The following are the possible values: • 10—fixed to 10Mbps • 100—fixed to 1000Mbps • 1000—fixed to 1000Mbps • auto—auto speed sensing (10/100/1000Mbps)
Default	auto	
Configuration Mode	config interface management	
History	3.1.0000	
Example	`switch (config interface mgmt0)# speed auto`	
Related Commands	show interfaces <ifname>	
Notes	• Setting the speed to “auto” also sets the duplex to “auto” • Setting the speed to one of the manual settings (generally “10”, “100”, or “1000”) also sets the duplex to a manual setting which is determined by querying the interface to find out its current auto-detected state	

4.1.8.1.8 dhcp

	dhcp [renew] no dhcp Enables DHCP on the specified interface. The no form of the command disables DHCP on the specified interface.	
Syntax Description	renew	Forces a renewal of the IP address. A restart on the DHCP client for the specified interface will be issued.
Default	Could be enabled or disabled (per part number) manufactured with 3.2.0500	
Configuration Mode	config interface management	
History	3.1.0000 3.9.1900: Added note	
Example	`switch (config interface mgmt0)# dhcp`	
Related Commands	show interfaces <ifname> configured	
Notes	• When enabling DHCP, the IP address and netmask are received via DHCP hence, the static IP address configuration is ignored • Enabling DHCP disables zeroconf and vice versa • Setting a static IP address and netmask does not disable DHCP. DHCP is disabled using the “no” form of this command, or by enabling zeroconf. • When static IP is configured, DHCP will not run.	

4.1.8.1.9 dhcp hostname

	dhcp hostname no dhcp hostname Enables fetching the hostname from DHCP for this interface. The no form of the command disables fetching the hostname from DHCP for this interface.
Syntax Description	N/A

Default	Enabled
Configuration Mode	config interface management
History	3.5.1000
Example	`switch (config interface mgmt0)# dhcp hostname`
Related Commands	hostname <hostname> show interfaces <ifname> configured
Notes	• If a hostname is configured manually by the user, that configuration would override the “dhcp hostname” configuration • When a default hostname is not configured, the DHCP server assigns the new hostname for your machine (after upgrading to version 3.5.1000) • These commands do not work on in-band interfaces

4.1.8.1.10 shutdown

	shutdown no shutdown Disables the specified interface. The no form of the command enables the specified interface.
Syntax Description	N/A
Default	no shutdown
Configuration Mode	config interface management
History	3.1.0000
Example	`switch (config interface mgmt0)# no shutdown`
Related Commands	show interfaces <ifname> configured
Notes	

4.1.8.1.11 zeroconf

	zeroconf no zeroconf Enables zeroconf on the specified interface. It randomly chooses a unique link-local IPv4 address from the 169.254.0.0/16 block. This command is an alternative to DHCP. The no form of the command disables the use of zeroconf on the specified interface.
Syntax Description	N/A
Default	no zeroconf
Configuration Mode	config interface management
History	3.1.0000
Example	`switch (config interface mgmt0)# zeroconf`
Related Commands	show interfaces <ifname> configured

Notes	Enabling zeroconf disables DHCP and vice versa.

4.1.8.1.12 comment

	comment <comment> no comment Adds a comment for an interface. The no form of the command removes a comment for an interface.	
Syntax Description	comment	A free-form string that has no semantics other than being displayed when the interface records are listed.
Default	no comment	
Configuration Mode	config interface management	
History	3.1.0000	
Example	`switch (config interface mgmt0)# comment my-interface`	
Related Commands		
Notes		

4.1.8.1.13 ipv6 enable

	ipv6 enable no ipv6 enable Enables all IPv6 addressing for this interface. The no form of the command disables all IPv6 addressing for this interface.	
Syntax Description	N/A	
Default	IPv6 addressing is disabled	
Configuration Mode	config interface management	
History	3.1.0000	
Example	`switch (config interface mgmt0)# ipv6 enable`	
Related Commands	ipv6 address show interface <ifname>	
Notes	• The interface identifier is a 64-bit long modified EUI-64, which is based on the MAC address of the interface • If IPv6 is enabled on an interface, the system will automatically add a link-local address to the interface. Link-local addresses can only be used to communicate with other hosts on the same link, and packets with link-local addresses are never forwarded by a router. • A link-local address, which may not be removed, is required for proper IPv6 operation. The link-local addresses start with “fe80::”, and are combined with the interface identifier to form the complete address.	

4.1.8.1.14 ipv6 address

	ipv6 address {<IPv6 address/netmask> | autoconfig [default | privacy]} no ipv6 {<IPv6 address/netmask> | autoconfig [default | privacy]} Configures IPv6 address and netmask to this interface, static or autoconfig options are possible. The no form of the command removes the given IPv6 address and netmask or disables the autoconfig options.	
Syntax Description	IPv6 address/ netmask	Configures a static IPv6 address and netmask. Format example: 2001:db8:1234::5678/64.
	autoconfig	Enables IPv6 stateless address auto configuration (SLAAC) for this interface. An address will be automatically added to the interface based on an IPv6 prefix learned from router advertisements, combined with an interface identifier.
	autoconfig default	Enables default learning routes. The default route will be discovered automatically, if the autoconfig is enabled.
	autoconfig privacy	Uses privacy extensions for SLAAC to construct the autoconfig address, if the autoconfig is enabled.
Default	No IP address available, auto config is enabled	
Configuration Mode	config interface management	
History	3.1.0000	
Example	`switch (config interface mgmt0)# ipv6 fe80::202:c9ff:fe5e:a5d8/64`	
Related Commands	ipv6 enable show interface <ifname>	
Notes	• On a given interface, up to 16 addresses can be configured • For Ethernet, the default interface identifier is a 64-bit long modified EUI-64, which is based on the MAC address of the interface	

4.1.8.1.15 ipv6 dhcp primary-intf

	ipv6 dhcp primary-intf <if-name> no ipv6 dhcp primary-intf Sets the interface from which non-interface-specific (resolver) configuration is accepted via DHCPv6. The no form of the command resets non-interface-specific (resolver) configuration.	
Syntax Description	if-name	Interface name: • lo • mgmt0 • mgmt1
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config)# ipv6 dhcp primary-intf mgmt0`	

Related Commands	ipv6 enable ipv6 address show interface <ifname>
Notes	

4.1.8.1.16 ipv6 dhcp stateless

	ipv6 dhcp stateless no ipv6 dhcp stateless Enables stateless DHCPv6 requests. The no form of the command disables stateless DHCPv6 requests.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	 3.1.0000
Example	`switch (config)# ipv6 dhcp stateless`
Related Commands	ipv6 enable ipv6 address show interface <ifname>
Notes	• This command only gets DNS configuration, not an IPv6 address • The no form of the command requests all information, including an IPv6 address

4.1.8.1.17 ipv6 dhcp client enable

	ipv6 dhcp client enable no ipv6 dhcp client enable Enables DHCPv6 on this interface. The no form of the command disables DHCPv6 on this interface.
Syntax Description	N/A
Default	ipv6 dhcp client enable
Configuration Mode	config interface management
History	 3.7.11xx 3.9.1900: Added note
Example	`switch (config interface mgmt0)# ipv6 dhcp client enable`
Related Commands	ipv6 dhcp client renew show ipv6 dhcp
Notes	When static IP is configured, DHCP will not run.

4.1.8.1.18 ipv6 dhcp client renew

	ipv6 dhcp client renew Renews DHCPv6 lease for this interface.

Syntax Description	N/A
Default	N/A
Configuration Mode	config interface management
History	3.7.11xx
Example	`switch (config interface mgmt0)# ipv6 dhcp client renew`
Related Commands	ipv6 dhcp client enable show ipv6 dhcp
Notes	

4.1.8.1.19 show interfaces mgmt0

	show interface mgmt0 Displays information on the management interface configuration and status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.6.8008: Updated example 3.9.1900: Updated example—added new output option of "no (Static IP is configured)"

Example	switch (config)# show interfaces mgmt0
	
	Interface mgmt0 status:
	  Comment :
	  Admin up : yes
	  Link up : yes
	  DHCP running : no (Static IP is configured)
	  IP address : 10.12.67.33
	  Netmask : 255.255.255.128
	  IPv6 enabled : yes
	  Autoconf enabled: no
	  Autoconf route : yes
	  Autoconf privacy: no
	  DHCPv6 running : no (Static IP is configured)
	  IPv6 addresses : 2
	
	IPv6 address:
	  1::1/64
	  fe80::7efe:90ff:fe65:dea8/64
	
	Speed : UNKNOWN
	  Duplex : full
	  Interface type : ethernet
	  Interface source: bridge
	  Bonding master : vrf_vrf-default
	  MTU : 1500
	  HW address : 7C:FE:90:65:DE:A8
	
	Rx:
	  13840892 bytes
	  58605 packets
	  0 mcast packets
	  2 discards
	  0 errors
	  0 overruns
	  0 frame
	
	Tx:
	  3796 bytes
	  38 packets
	  0 discards
	  0 errors
	  0 overruns
	  0 carrier
	  0 collisions
	  1000 queue len

Related Commands	
Notes	

4.1.8.1.20 show interfaces mgmt0 brief

	show interface mgmt0 brief Displays brief information on the management interface configuration and status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.6.8008: Updated example

Example	switch (config)# show interfaces mgmt0 brief
	
	Interface mgmt0 status:
	  Comment :
	  Admin up : yes
	  Link up : yes
	  DHCP running : yes
	  IP address : 10.12.67.33
	  Netmask : 255.255.255.128
	  IPv6 enabled : yes
	  Autoconf enabled: no
	  Autoconf route : yes
	  Autoconf privacy: no
	  DHCPv6 running : yes (but no valid lease)
	  IPv6 addresses : 1
	
	  IPv6 address:
	  fe80::268a:7ff:fe53:3d8e/64
	
	  Speed : 1000Mb/s (auto)
	  Duplex : full (auto)
	  Interface type : ethernet
	  Interface source: bridge
	  MTU : 1500
	  HW address : 24:8a:07:53:3d:8e

Related Commands	
Notes	

4.1.8.1.21 show interfaces mgmt0 configured

	show interface mgmt0 configured Displays configuration information about the specified interface.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	 3.1.0000 3.5.1000: Updated example with “DHCP Hostname” 3.6.8008: Updated example

Example	switch (config)# show interfaces mgmt0 configured
	
	Interface mgmt0 configuration:
	  Comment :
	  Enabled : yes
	  DHCP : yes
	  DHCP Hostname : yes
	  Zeroconf : no
	  IP address :
	  Netmask :
	  IPv6 enabled : yes
	  Autoconf enabled: no
	  Autoconf route : yes
	  Autoconf privacy: no
	  DHCPv6 enabled : yes
	  IPv6 addresses : 0
	  Speed : auto
	  Duplex : auto
	  MTU : 1500

Related Commands	
Notes	

4.1.8.2 Hostname Resolution

4.1.8.2.1 hostname

	hostname <hostname> no hostname Sets a static system hostname. The no form of the command clears the system hostname.	
Syntax Description	hostname	A free-form string
Default	Default hostname	
Configuration Mode	config	
History	3.1.0000 3.6.3004: Added support for the character “.”	
Example	`switch (config)# hostname my-switch-hostname`	
Related Commands	show hosts	
Notes	• Hostname may contain letters, numbers, periods (‘.’), and hyphens (‘-’), in any combination • Hostname may be 1-63 characters long • Hostname may not begin with a hyphen • Hostname may not contain other characters, such as “%”, “_” etc. • Hostname may not be set to one of the valid logging commands (i.e. debug-files, fields, files, format, level, local, monitor, receive, trap) • Changing the hostname stamps a new HTTPS certificate	

4.1.8.2.2 ip name-server

	ip name-server <IPv4/IPv6 address> no ip name-server <IPv4/IPv6 address> Sets the static name server. The no form of the command clears the name server.	
Syntax Description	IPv4/IPv6 address	IPv4 or IPv6 address.
Default	No server name	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config)# ip name-server 9.9.9.9`	
Related Commands	show hosts	
Notes		

4.1.8.2.3 ip domain-list

	ip domain-list <domain-name> no ip domain-list <domain-name> Sets the static domain name. The no form of the command clears the domain name.	
Syntax Description	domain-name	The domain name in a string form. A domain name is an identification string that defines a realm of administrative autonomy, authority, or control in the Internet.
Default	No static domain name	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config)# ip domain-list mydomain.com`	
Related Commands	show hosts	
Notes		

4.1.8.2.4 ip/ipv6 host

	{ip | ipv6} host <hostname> <ip-address> no {ip | ipv6} host <hostname> <ip-address> Configures the static hostname IPv4 or IPv6 address mappings. The no form of the command clears the static mapping.	
Syntax Description	hostname	The hostname in a string form.
	IP Address	The IPv4 or IPv6 address.
Default	No static domain name	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config)# ip host my-host 2.2.2.2` `switch (config)# ipv6 host my-ipv6-host 2001::8f9`	
Related Commands	show hosts	
Notes		

4.1.8.2.5 ip/ipv6 map-hostname

	{ip |ipv6} map-hostname no {ip | ipv6} map-hostname Maps between the currently-configured hostname and the loopback address 127.0.0.1. The no form of the command clears the mapping.
Syntax Description	N/A
Default	IPv4 mapping is enabled by default IPv6 mapping is disabled by default

Configuration Mode	config
History	3.1.0000
Example	`switch (config)# ip map-hostname`
Related Commands	show hosts
Notes	• If no mapping is configured, a mapping between the hostname and the IPv4 loopback address 127.0.0.1 will be added • The no form of the command maps the hostname to the IPv6 loopback address if there is no statically configured mapping from the hostname to an IPv6 address (disabled by default) • Static host mappings are preferred over DNS results. As a result, with this option set, you will not be able to look up your hostname on your configured DNS server; but without it set, some problems may arise if your hostname cannot be looked up in DNS.

4.1.8.2.6 show hosts

	show hosts Displays hostname, DNS configuration, and static host mappings.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.8.1000: Updated example

Example	`switch (config)# show hosts`
	
	`Hostname: switch1`
	
	`Name servers:`
	` 10.7.77.192 dynamic (DHCP on mgmt0)`
	` 10.7.77.135 dynamic (DHCP on mgmt0)`
	` 10.198.0.169 dynamic (DHCP on mgmt0)`
	` (*) 10.211.0.124 dynamic (DHCP on mgmt0)`
	
	`Domain names:`
	` mtl.labs.mlnx dynamic (DHCP on mgmt0)`
	
	`(*) Inactive due to system limits on name servers and domain names.`
	
	`Static IPv4 host mappings:`
	` 10.7.144.133 --> switch1`
	` 127.0.0.1 --> localhost`
	
	`Static IPv6 host mappings:`
	` ::1 --> localhost6`
	
	`Automatically map hostname to loopback address : yes`
	`Automatically map hostname to IPv6 loopback address: no`

Related Commands	
Notes	

4.1.8.3 Routing

4.1.8.3.1 IP route

	{ip | ipv6} route [vrf <vrf-name>] {<network-prefix> <netmask> | <network-prefix>/<masklen>} <next-hop> no ip route [vrf <vrf-name>] {<network-prefix> <netmask> | <network-prefix>/<masklen>} <next-hop> Sets a static route for a given IP. The no form of the command deletes the static route.	
Syntax Description	network-prefix	IPv4 or IPv6 network prefix
	netmask	IPv4 netmask formats are: • /24 • 255.255.255.0 IPv6 netmask format is: • /48 (as a part of the network prefix)
	nexthop-address	The IPv4 or IPv6 address of the next hop router for this route
	ifname	The interface name (e.g., mgmt0, mgmt1)
	vrf-name—VRF session name	
Default	N/A	
Configuration Mode	config	
History	3.1.0000 3.9.2000—Added VRF option	
Example	`switch (config)# ip route 20.20.20.0 255.255.255.0 mgmt0`	
Related Commands	show ip route	
Notes		

4.1.8.3.2 ipv6 default-gateway

	ipv6 default-gateway {<ip-address> | <ifname>} no ipv6 default-gateway Sets a static default gateway. The no form of the command deletes the default gateway.	
Syntax Description	ip address	The default gateway IP address (IPv6)
	ifname	The interface name (e.g.,, mgmt0, mgmt1)
Default	N/A	
Configuration Mode	config	
History	3.1.0000 3.2.0500: Removed IPv4 configuration option	
Example	`switch (config)# ipv6 default-gateway ::1`	
Related Commands	show ip/ipv6 route show ipv6 default-gateway	

Notes	• The configured default gateway will not be used if DHCP is enabled • In order to configure ipv4 default-gateway use ‘ip route’ command.

4.1.8.3.3 show ip/ipv6 route

	show {ip | ipv6} route [static] Displays the routing table in the system.	
Syntax Description	static	Filters the table with the static route entries
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	

Example

```
switch (config)# show ip route
Destination       Mask                Gateway         Interface       Source
default           0.0.0.0             172.30.0.1      mgmt0           DHCP
10.10.10.10       255.255.255.255     0.0.0.0         mgmt0           static
20.10.10.10       255.255.255.255     172.30.0.1      mgmt0           static
20.20.20.0        255.255.255.0       0.0.0.0         mgmt0           static
172.30.0.0        255.255.0.0         0.0.0.0         mgmt0           interface
switch (config)# show ipv6 route
Destination prefix
Gateway                                 Interface      Source
-------------------------------------------------------------------------
::/0
    ::                                  mgmt0          static
::1/128
    ::                                  lo             local
2222:2222:2222::/64
    ::                                  mgmt1          interface
```

Related Commands	ip route
Notes	

4.1.8.3.4 show ipv6 default-gateway

	show ipv6 default-gateway [static] Displays the default gateway.	
Syntax Description	static	Displays the static configuration of the default gateway
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	

Example	`switch (config)# show ipv6 default-gateway`
	`Active default gateways:`
	`  172.30.0.1 (interface: mgmt0)`
	`switch (config)# show ipv6 default-gateway static`
	`Configured default gateway: 10.10.10.10`	

Related Commands	ipv6 default-gateway	
Notes	The configured IPv4 default gateway will not be used if DHCP is enable	

4.1.8.4 Network to Media Resolution (ARP & NDP)

IPv4 network use Address Resolution Protocol (ARP) to resolve IP address to MAC address, while IPv6 network uses Network Discovery Protocol (NDP) that performs basically the same as ARP.

4.1.8.4.1 ip arp

	ip arp <ip-address> <mac-address> no ip arp <ip-address> <mac-address> Sets a static ARP entry. The no form of the command deletes the static ARP.	
Syntax Description	ip-address	IPv4 address
	mac-address	MAC address
Default	N/A	
Configuration Mode	config interface management	
History	3.2.0500	
Example	`switch (config interface mgmt0)#ip arp 20.20.20.20 aa:aa:aa:aa:aa:aa`	
Related Commands	show ip arp ip route	
Notes		

4.1.8.4.2 ip arp timeout

	ip arp [vrf <vrf-name>] time out <timeout-value> no ip arp [vrf <vrf-name>] timeout Sets the dynamic ARP cache timeout. The no form of the command sets the timeout to default.	
Syntax Description	timeout-value	Time (in seconds) that an entry remains in the ARP cache Range: 60-28800
	vrf-name	VRF session name
Default	1500 seconds	
Configuration Mode	config	
History	3.2.0230	
	3.5.1000	Added VRF parameter and updated Notes
Example	`switch (config) # ip arp timeout 2000`	
Related Commands	ip arp show ip arp	
Notes	• This value is used as the default ARP timeout whenever a new IP interface is created • The time interval after which each ARP entry becomes stale may actually vary from 50-150% of the configured value	

4.1.8.4.3 ip arp cache-update

	ip arp cache-update {always | garp} no ip arp cache-update {always | garp} Sets up the cache update strategy. The no form of the command sets cache to default.	
Syntax Description	always	Any broadcast ARP will cause a neighbor cache update
	garp	Only GARP messages will cause a cache update
Default	garp	
Configuration Mode	config interface port-channel config interface ethernet config interface vlan	
History	3.9.3100	
Example	`switch (config interface ethernet 1/1) # ip arp cache-update always`	
Related Commands		
Notes		

4.1.8.4.4 show ip arp

	show ip arp [interface <type> | <ip-address> | count] Displays ARP table.	
Syntax Description	interface type	Filters the table according to a specific interface (i.e. mgmt0)
	ip-address	Filters the table to the specific ip-address
	count	Shows ARP statistics
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.3000	
Example		

```
switch (config)# show ip arp

Total number of entries: 3

  Address              Type           Hardware Address     Interface
  ------------------------------------------------------------------------
  10.209.0.1           Dynamic ETH    00:00:5E:00:01:01    mgmt0
  10.209.1.120         Dynamic ETH    00:02:C9:62:E8:C2    mgmt0
  10.209.1.121         Dynamic ETH    00:02:C9:62:E7:42    mgmt0
switch (config)# show ip arp count
ARP Table size: 3 (inband: 0, out of band: 3)
```

Related Commands	
Notes	

4.1.8.4.5 ipv6 neighbor

	ipv6 neighbor <ipv6-address> <ifname> <mac-address> no ipv6 neighbor <ipv6-address> <ifname> <mac-address> Adds a static neighbor entry. The no form of the command deletes the static entry.	
Syntax Description	ipv6-address	The IPv6 address
	ifname	The management interface (i.e. mgmt0, mgmt1)
	mac-address	The MAC address
Default	N/A	
Configuration Mode	config	
History	 3.1.0000	
Example	`switch (config)# ipv6 neighbor 2001:db8:701f::8f9 mgmt0 00:11:22:33:44:55`	
Related Commands	show ipv6 neighbor ipv6 route arp clear ipv6 neighbors	
Notes	• ARP is used only with IPv4. In IPv6 networks, Neighbor Discovery Protocol (NDP) is used similarly. • Use The no form of the command to remove static entries. Dynamic entries can be cleared via the “clear ipv6 neighbors” command.	

4.1.8.4.6 clear ipv6 neighbors

	clear ipv6 neighbors{ethernet <port> | vlan <vlan-id> | port-channel <id> | vrf <vrf-id>} [<ip-addr>] Clears the dynamic neighbors cache.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	 3.1.0000 3.6.4110: Updated command
Example	`switch (config)# clear ipv6 neighbors`
Related Commands	ipv6 neighbor show ipv6 neighbor arp
Notes	• Clearing Neighbor Discovery Protocol (NDP) cache removes only the dynamic entries learned and not the static entries configured • Use the no form of the command to remove static entries

4.1.8.4.7 show ipv6 neighbors

	show ipv6 neighbors [static] Displays the Neighbor Discovery Protocol (NDP) table.	
Syntax Description	static	Filters only the table of the static entries.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example		

switch (config)# show ipv6 neighbors
`IPv6 Address Age MAC Address State Interf`
`---------- ----- ----------- ------- ------`
`2001::2 9428 aa:aa:aa:aa:aa:aa permanent mgmt0`		

Related Commands	ipv6 neighbor clear ipv6 neighbor show ipv6	
Notes		

4.1.8.5 DHCP

4.1.8.5.1 ip dhcp

	ip dhcp {default-gateway yield-to-static | hostname <hostname>| primary-intf <ifname> | send-hostname} no ip dhcp {default-gateway yield-to-static| hostname | | primary-intf | send-hostname} Sets global DHCP configuration. The no form of the command deletes the DHCP configuration.	
Syntax Description	yield-to-static	Does not allow you to install a default gateway from DHCP if there is already a statically configured one.
	hostname	Specifies the hostname to be sent during DHCP client negotiation if send-hostname is enabled.
	primary-intf <ifname>	Sets the interface from which a non-interface-specific configuration (resolver and routes) will be accepted via DHCP. Default: "primary-intf mgmt0"
	send-hostname	Enables the DHCP client to send a hostname during negotiation.
Default	Disabled	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config)# ip dhcp default-gateway yield-to-static`	
Related Commands	show ip dhcp dhcp [renew]	

Notes	DHCP is supported for IPv4 networks only

4.1.8.5.2 show ip dhcp

	show ip dhcp Displays the DHCP configuration and status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	 3.1.0000 3.6.5000: Updated example

Example	`switch (config)# show ip dhcp`
	`----------`
	`Interface DHCP DHCP Valid`
	`Enabled Running lease`
	`----------`
	`dummy0 no no no`
	`lo no no no`
	`mgmt0 yes yes yes`
	`mgmt1 no no no`
	`mgmts0 no no no`
	`mgmts1 no no no`
	`vif1 no no no`
	
	`IPv4 dhcp default gateway yields to static configuration: no`
	
	`DHCP primary interface:`
	`Configured: mgmt0`
	`Active: mgmt0`
	
	`DHCP client options:`
	`Send Hostname: no`
	`Client Hostname: 1.1.1.1`

Related Commands	ip dhcp dhcp [renew]
Notes	

4.1.8.6 General IPv6

4.1.8.6.1 ipv6 enable

	ipv6 enable no ipv6 enable Enables IPv6 globally on the management interface. The no form of the command disables IPv6 globally on the management interface.
Syntax Description	N/A
Default	IPv6 is disabled
Configuration Mode	config
History	3.1.0000
Example	`switch (config)# ipv6 enable`

Related Commands	ipv6 default-gateway ipv6 host ipv6 map-hostname ipv6 neighbor ipv6 route show ipv6 show ipv6 default-gateway show ipv6 route
Notes	

4.1.8.7 IP Diagnostic Tools

4.1.8.7.1 ping

	ping [-LRUbdfnqrvVaA] [-c count] [-i interval] [-w deadline] [-p pattern] [-s packetsize] [-t ttl] [-I interface or address] [-M mtu discovery hint] [-S sndbuf] [-T timestamp option] [-Q tos] [hop1 ...] destination Sends ICMP echo requests to a specified host.	
Syntax Description	Linux Ping options	https://www.lifewire.com/uses-of-command-ping-2201076
Default	N/A	
Configuration Mode	config	
History	3.1.0000	

Example	`switch (config)# ping 172.30.2.2`
	`PING 172.30.2.2 (172.30.2.2) 56(84) bytes of data.`
	`64 bytes from 172.30.2.2: icmp_seq=1 ttl=64 time=0.703 ms`
	`64 bytes from 172.30.2.2: icmp_seq=2 ttl=64 time=0.187 ms`
	`64 bytes from 172.30.2.2: icmp_seq=3 ttl=64 time=0.166 ms`
	`64 bytes from 172.30.2.2: icmp_seq=4 ttl=64 time=0.161 ms`
	`64 bytes from 172.30.2.2: icmp_seq=5 ttl=64 time=0.153 ms`
	`64 bytes from 172.30.2.2: icmp_seq=6 ttl=64 time=0.144 ms`
	`...`
	`--- 172.30.2.2 ping statistics ---`
	`6 packets transmitted, 6 received, 0% packet loss, time 5004ms`
	`rtt min/avg/max/mdev = 0.144/0.252/0.703/0.202 ms`	

Related Commands	traceroutes	
Notes		

4.1.8.7.2 traceroute

	traceroute [-46dFITUnrAV] [-f first_ttl] [-g gate,...] [-i device] [-m max_ttl] [-N squeries] [-p port] [-t tos] [-l flow_label] [-w waittime] [-q nqueries] [-s src_addr] [-z sendwait] host [packetlen] Traces the route packets take to a destination.	
Syntax Description	-4	Uses IPv4
	-6	Uses IPv6
	-d	Enables socket level debugging
	-F	Sets DF (do not fragment bit) on
	-I	Uses ICMP ECHO for tracerouting

	-T	Uses TCP SYN for tracerouting
	-U	Uses UDP datagram (default) for tracerouting
	-n	Does not resolve IP addresses to their domain names
	-r	Bypasses the normal routing and send directly to a host on an attached network
	-A	Performs AS path lookups in routing registries and print results directly after the corresponding addresses
	-V	Prints version info and exit
	-f	Starts from the first_ttl hop (instead from 1)
	-g	Routes packets through the specified gateway (maximum 8 for IPv4 and 127 for IPv6)
	-i	Specifies a network interface with which to operate
	-m	Sets the max number of hops (max TTL to be reached). Default is 30.
	-N	Sets the number of probes to be tried simultaneously (default is 16)
	-p	Uses destination port. It is an initial value for the UDP destination port (incremented by each probe, default is 33434), for the ICMP seq number (incremented as well, default from 1), and the constant destination port for TCP tries (default is 80).
	-t	Sets the TOS (IPv4 type of service) or TC (IPv6 traffic class) value for outgoing packets
	-l	Uses specified flow_label for IPv6 packets
	-w	Sets the number of seconds to wait for response to a probe (default is 5.0). Non-integer (float point) values allowed too.
	-s	Uses source src_addr for outgoing packets.
	-q	Sets the number of probes per each hop. Default is 3.
	-z	Sets minimal time interval between probes (default is 0). If the value is more than 10, then it specifies a number in milliseconds, else it is a number of seconds (float point values allowed too).
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example		
switch (config)# traceroute 192.168.10.70 traceroute to 192.168.10.70 (192.168.10.70), 30 hops max, 40 byte packets 1 172.30.0.1 (172.30.0.1) 3.632 ms 2.849 ms 3.544 ms 2 10.222.128.46 (10.222.128.46) 3.176 ms 3.289 ms 3.656 ms 3 10.158.128.30 (10.158.128.30) 15.331 ms 15.819 ms 16.388 ms 4 10.158.128.65 (10.158.128.65) 20.468 ms 7.893 ms 12.27 ms 5 10.7.34.115 (10.7.34.115) 16.405 ms 11.985 ms 12.264 ms6 192.168.10.70 (192.168.10.70) 16.377 ms 16.091 ms 20.475 ms		
Related Commands	ping	
Notes		

4.1.8.7.3 tcpdump

	tcpdump [-aAdDeflLnNOpqRStuUvxX] [-c count] [-C file_size] [-E algo:secret] [-F file] [-i interface] [-M secret] [-r file] [-s snaplen] [-T type] [-w file] [-W filecount] [-y datalinktype] [-Z user] [-D list possible interfaces] [expression] Invokes standard binary, passing command line parameters straight through. Runs in foreground, printing packets as they arrive, until the user hits Ctrl+C.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Example	switch (config)# tcpdump ..... 09:37:38.678812 IP 192.168.10.7.ssh > 192.168.10.1.54155: P 1494624:1494800(176) ack 625 win 90 <nop,nop,timestamp 5842763 858672398> 09:37:38.678860 IP 192.168.10.7.ssh > 192.168.10.1.54155: P 1494800:1495104(304) ack 625 win 90 <nop,nop,timestamp 5842763 858672398> ... 9141 packets captured 9142 packets received by filter 0 packets dropped by kernel
Related Commands	
Notes	

4.1.9 Control Plane Policing (CoPP)

Control Plane Policing or Policies (CoPP) ensures the CPU and control plane are not over-utilized which is essential for the robustness of the switch. CoPP limits the number of control plane packets.

This software implements several CoPP mechanisms:

- ACLs may be used to limit the rate of packets or bytes of a certain type, including L3 control packets (L2 control packets are forwarded to the CPU before the ACL)
- Policers on traffic going to the CPU—these policers are configured by the operating system and cannot be modified by the user
- IP filter tables limit the traffic to the CPU coming in from the management ports

4.1.9.1 IP Table Filtering

IP table filtering is a mechanism that allows the user to apply actions to a specific control packet flow identified by a certain flow key.

This mechanism is used in order to protect switch control traffic against attacks. For example, it could allow traffic coming from a specific trusted management subnet only, block the SNMP UDP port from receiving traffic, and force ping rate to be lower than a specific threshold.
Each IP table rule is defined by key, priority, and action:

- Key—the key is a combination of physical port and layer 3 parameters (e.g. SIP, DIP, SPORT, DPORT, etc.), and other fields. Each part of the key, can be set to a specific value or masked.
- Priority—each rule in the IP table is assigned a priority, and the rule with the highest priority whose key matches the packet executes the action.
- Action—the action describes the behavior of packets which match the key. The action type may be drop, accept, rate limit, etc.

An IP-table rule is bound to an IP interface that can be a management out-of-band interface, VLAN interface, or router port interface. Once bound, all traffic received (ingress rule) or transmitted (egress rule) in this direction is being verified with all bounded rules.

Once a match was found, the rule action is executed. If no match is found, the default policy of the chain shall apply.

> IP table rules get a lower priority than ACL mechanism.

> In the rare case that IP filter is used while the input policy is "drop" (i.e., ip filter chain input policy drop) and an NTP server or an MLAG switch is used, then the following rule needs to be added that allows src-ip 127.0.0.1 (which is a requirement for any clustered application (e.g., mlag-vip) and NTP):
> `ip filter chain input rule append tail target accept dup-delete source-addr 127.0.0.1 /32`

4.1.9.1.1 Configuring IP Table Filtering

Prerequisite for IPv6:

```
switch (config) # ipv6 enable
```

To configure IPv4 table filtering:

1. Select the policy that applies to the input/output chain (default is “accept”). Run:

   ```
   switch (config)# ip filter chain input policy drop
   switch (config)# ip filter chain output policy accept
   ```

2. Append filtering rules to the list or set a specific rule number, select a target, and (optional) any additional filter conditions. For example, run:

   ```
   switch (config) # ip filter chain input rule append tail target rate-limit 2 protocol udp
   switch (config) # ip filter chain input rule set 2 target drop protocol icmp in-intf mgmt1
   switch (config) # ip filter chain output rule append tail target drop protocol icmp
   ```

3. Enable IP table filtering. Run:

   ```
   switch (config) # ip filter enable
   ```

4. Verify IP table filtering configuration. Run:

   ```
   switch (config) # show ip filter configured
   ```

```
Packet filtering for IPv4: enabled

IPv4 configuration:
  Chain 'input' Policy 'accept':
    Rule 1:
      Target      : rate-limit 2 pps
      Protocol    : udp
      Source      : all
      Destination : all
      Interface   : all
      State       : any
      Other Filter:  -

    Rule 2:
      Target      : drop
      Protocol    : icmp
      Source      : all
      Destination : all
      Interface   : mgmt1 (ingress)
      State       : any
      Other Filter:  -

  Chain 'output' Policy 'accept':
    Rule 1:
      Target      : drop
      Protocol    : icmp
      Source      : all
      Destination : all
      Interface   : all
      State       : any
      Other Filter: -
```

4.1.9.1.2 Modifying IP Table Filtering

To modify IP table filtering configuration:

```
switch (config) # ip filter chain input rule modify 3 target reject-with icmp6-adm-prohibited source-addr 10::0 /
126
```

To delete an existing IP table filtering rule:

```
switch (config) # no ip filter chain input rule 2
```

To delete all existing IP table filtering rules:

```
switch (config) # no ip filter chain output rule all
```

To insert an IP table filtering rule in a chain:

```
switch (config) # ip filter chain input rule 2 set target drop protocol tcp dest-port 22 in-intf mgmt1
```

4.1.9.1.3 Rate-Limit Rule Configuration

Using a rate-limit target allows to create a rule to limit the rate of certain traffic types. The limit is specified in packets per second (pps) and can be anywhere between 1-1000 pps. When enabled, the system takes the user specified rate and converts it into units of 1/10000 of a second. Therefore, any value greater than 100 can have a slight difference when the rule is displayed using the show command.

Unlike other rules which are a match type of rule, limiting packets should be followed by a rule that drops additional packets of the same “type”. Alternatively, this can be implicitly achieved by setting

the chain policy to “drop” so that it drops packets not processed by matching rules. Otherwise, no effect of the rule is observed as the remaining traffic simply gets accepted.

> Rate-limit is implemented with an average rate and a burst-limit. Rate values are specified in pps and take a range from 1-1000 pps. For rate values in the range 1-100, the burst value is set equal to the rate value. For rate values in the range 101-1000, the burst limit is set to 100.

4.1.9.1.4 IP Table Filtering Default Rules

IP table filtering is enabled and Firewall default IP filter rules are applied.

- To reset/apply default rules on system, run the command “ip filter reset-to-default-rules”
- To enable IP Filter, run the command “ip filter enable”
- To list the default firewall rules, run the command “show ip filter”
- Note when touching a default rule (delete/move/modify) all IP Filter rules will be reflected on “show running-config”, to restore default rules, run the command “ip filter reset-to-default-rules”
- Restoring factory default configuration will reset the default rules and enable the feature

4.1.9.1.4.1 Firewall Default Rules

Prerouting-Mangle Chain Rules
• ip filter chain prerouting-mangle rule append tail target drop in-intf mgmt0 protocol tcp conntrack new tcp-op-mss mss-not-in-range 536:65535 not-dest-port 22
Input Chain Rules

- ip filter chain input rule append tail target accept in-intf lo
- ip filter chain input rule append tail target drop in-intf mgmt0 dest-addr 127.0.0.0 /8
- ip filter chain input rule append tail target accept in-intf mgmt0 state established,related
- ip filter chain input rule append tail target drop in-intf mgmt0 protocol tcp tcp-op syn match-not-syn state new
- ip filter chain input rule append tail target drop in-intf mgmt0 fragment enable
- ip filter chain input rule append tail target drop in-intf mgmt0 protocol tcp tcp-op flags all
- ip filter chain input rule append tail target drop in-intf mgmt0 protocol tcp tcp-op flags none
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol tcp tcp-op flags reset rate-limit 2/second burst-limit 2",
- ip filter chain input rule append tail target drop in-intf mgmt0 state invalid
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol tcp state new rate-limit 50/second burst-limit 50
- ip filter chain input rule append tail target drop in-intf mgmt0 protocol tcp state new tcp-op syn match-not-syn not-dest-port 22
- ip filter chain input rule append tail target drop in-intf mgmt0 recent name "portscan" recent rcheck-sec 86400
- ip filter chain input rule append tail target none in-intf mgmt0 recent name portscan recent remove
- ip filter chain input rule append tail target none in-intf mgmt0 protocol tcp dest-port 22 conntrack new recent set
- ip filter chain input rule append tail target drop in-intf mgmt0 protocol tcp dest-port 22 conntrack new recent update-sec 60 recent hitcount 10
- ip filter chain input rule append tail target none in-intf mgmt0 protocol tcp dest-port 443 conntrack new recent set
- ip filter chain input rule append tail target drop in-intf mgmt0 protocol tcp dest-port 443 conntrack new recent update-sec 60 recent hitcount 10
- ip filter chain input rule append tail target none in-intf mgmt0 protocol udp dest-port 161 conntrack new recent set
- ip filter chain input rule append tail target drop in-intf mgmt0 protocol udp dest-port 161 conntrack new recent update-sec 60 recent hitcount 25
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol tcp conntrack new rate-limit 60/second burst-limit 20
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol tcp dest-port 22 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol tcp dest-port 443 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol tcp dest-port 179 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol udp dest-port 68 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol udp dest-port 122 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol udp dest-port 161 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol udp dest-port 6306 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol udp dest-port 69 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol udp dest-port 389 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol tcp dest-port 389 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol udp dest-port 1812-1813 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol udp dest-port 49 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol tcp dest-port 49 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol udp source-port 53 conntrack new,established

- ip filter chain input rule append tail target accept in-intf mgmt0 protocol tcp source-port 53 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol udp dest-port 500 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol udp dest-port 4500 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol udp dest-port 1293 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol tcp dest-port 1293 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol udp dest-port 1707 conntrack new,established
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol tcp dest-port 1707 conntrack new,established
- ip filter chain input rule append tail target accept protocol udp dest-port 3786 conntrack new,established in-intf lo
- ip filter chain input rule append tail target accept protocol udp dest-port 33000 conntrack new,established in-intf lo
- ip filter chain input rule append tail target accept in-intf mgmt0 protocol icmp
- ip filter chain input rule append tail target accept dup-delete source-port 5353 dest-port 5353 in-intf mgmt0 protocol udp conntrack new,established
- ip filter chain input rule append tail target logging in-intf mgmt0

Output Chain Rules

- ip filter chain output rule append tail target drop out-intf mgmt0 state invalid
- ip filter chain output rule append tail target accept out-intf mgmt0

Logging Chain Rules

- ip filter chain logging rule append tail target nflog in-intf mgmt0 rate-limit 1/minute logging-options prefix "IPTables-Dropped-<Domain>: " logging-options group 3
- ip filter chain logging rule append tail target drop in-intf mgmt0

4.1.9.2 Control Plane Policing Commands

4.1.9.2.1 ip filter enable | ipv6 filter enable

	{ip | ipv6} filter enable no {ip | ipv6} filter enable Enables IP filtering. The no form of the command disables IP filtering.
Syntax Description	N/A
Default	ip Enabled ip6 Disabled
Configuration Mode	config
History	3.5.1000 3.10.3000 IP Filter is enabled by default
Example	`switch (config) # ip filter enable`
Related Commands	
Notes	It is recommended to run this command only after configuring all of the IP table filter parameters.

4.1.9.2.2 ip filter chain policy | ipv6 filter chain policy

	{ip | ipv6} filter chain <chain_name> policy {accept | drop} no {ip | ipv6} filter chain <chain_name> policy Configures default policy for a specific chain (if no rule matches this default policy action shall apply). The no form of the command resets default policy for a specific chain.	
Syntax Description	chain_name	Selects a chain for which to add or modify a filter: • input - input chain or ingress interfaces • output - output chain or egress interfaces
	accept	Accepts all traffic by default for this chain
	drop	Drops all traffic by default for this chain
Default	Accept for input and output chains	
Configuration Mode	config	
History	3.5.1000	
Example	`switch (config) # ipv6 filter chain input policy accept`	
Related Commands		
Notes		

4.1.9.2.3 ip filter chain rule target | ipv6 filter chain rule target

	{ip | ipv6} filter chain <chain_name> rule <oper> target <target> [<param>] no {ip | ipv6} filter chain <chain_name> rule {<number> | all} Inserts rule before specified rule number. The no form of the command deletes rule for a specific chain.	
Syntax Description	chain_name	A chain to which to add or modify a filter: • input - input chain or ingress interfaces • output - output chain or egress interfaces
	rule	• append tail - appends operation to the bottom of operation list • insert <oper_num> - inserts operation at specified position (existing operation at that position moves back in the list) • modify <oper_num> - modifies existing operation at specified position. Only the parameters specified in this invocation are altered; everything else is left untouched. • move <oper_num1> to <oper_num2> - moves one operation to another place in the operation list • set <oper_num> - sets operation at specified position (overwrites existing)
	target	• accept - allows the packets that match the rule into the management plane • drop - drops packets that match the rule • rate-limit - allows with rate limiting in packets per sec (PPS) • reject-with - drops the packet and replies with an ICMP error message

	param	• comment <text> - specifies description string for this rule (60 chars max) • dest-addr <ip> - IP matching a specific destination address or address range. A specific IPv4 address can be provided or an entire subnet by giving an address along with netmask in dot notation or as a CIDR notation (e.g. /24). • not-dest-addr <ip> - IP not matching a specific destination address range • dest-port <port(s)> - matching a specific destination port or port range • not-dest-port <port(s)> - port not matching a specific destination port or port range • dup-delete - deletes any preexisting duplicates of this rule • in-intf - interface matching a specific inbound interface • not-in-intf <if_name> - interface not matching a specific inbound interface • out-intf <if_name> - matches a specific outbound interface • not-out-intf <if_name> - interface not matching a specific outbound interface
	param4 (cont.)	• protocol <if_name> - matches a specific protocol • tcp • udp • icmp • all • not-protocol <protocol> - does not match a specific protocol • tcp • udp • icmp • all • source-addr <ip> - matches a specific source address range • not-source-addr <ip> - does not match a specific source address range • source-port <port(s)> - matches a specific source port or port range • not-source-port <port(s)> - does not match a specific source port or port range • state - matches packets in a particular state. Possible values: • established - packet associated with an established connection which has seen traffic in both directions • related - packet that starts a new connection but is related to an existing connection • new - packet that starts a new, unrelated connection • A combination can be entered separated by commas
Default	N/A	
Configuration Mode	config	
History	3.5.1000	
Example	`switch (config) # ipv6 filter enable chain input rule append tail target drop state related protocol all dup-delete`	
Related Commands		

Notes	• The source and destination ports may each be either a single number, or a range specified as “<low>-<high>”. For example: “10-20” would specify ports 10 through 20 (inclusive). • The port parameter only works in conjunction with TCP and UDP • Setting a “positive” rule removes any corresponding “not-” rules, and vice-versa • The “state” parameter is a classification of the packet relative to existing connections • If TCP or UDP are selected for the “protocol” parameter, source and/or destination ports may be specified. If ICMP is selected, these options are either ignored, or an error is produced.

4.1.9.2.4 ip filter options include-bridges

	{ip | ipv6} filter options include-bridges no {ip | ipv6} filter options include-bridges Applies IP filters to bridges
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.5.1000
Example	`switch (config) # ip filter options include-bridges`
Related Commands	
Notes	

4.1.9.2.5 ip filter reset-to-default-rules

	ip filter reset-to-default-rules Deletes all configured IP filter rules and add the default rules defined in the user manual under section "IP Table Filtering Default Rules", above.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.10.3000
Example	`switch (config) # ip filter reset-to-default-rules`
Related Commands	
Notes	

4.1.9.2.6 show ip filter

	show ip filter Displays IPv4 filtering state.
Syntax Description	N/A
Default	N/A

Configuration Mode	config
History	3.6.6000

Example	switch (config) # show ip filter
	
	Packet filtering for IPv4: enabled
	
	Active IPv4 filtering rules (omitting any not from configuration):
	  Chain 'input' Policy 'accept':
	  Rule 1:
	  Target : accept
	  Protocol : all
	  Source : all
	  Destination : 1.1.1.0/24
	  Interface : all
	  State : any
	  Other Filter: -
	
	Chain 'output' Policy 'accept':
	  Rule 1:
	  Target : reject-with icmp-net-unreachable
	  Protocol : tcp
	  Source : all
	  Destination : all
	  Interface : all
	  State : any
	  Other Filter: dest-port 1000

Related Commands	
Notes	

4.1.9.2.7 show ip filter all

	show ip filter all Displays IPv4 filtering state (including un-configured rules).
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000
Example	Destination : 1.1.1.0/24 Interface : all State : any Other Filter: - Chain 'output' Policy 'accept': Rule 1: Target : reject-with icmp-net-unreachable Protocol : tcp Source : all Destination : all Interface : all State : any Other Filter: dest-port 1000
Related Commands	
Notes	

4.1.9.2.8 show ip filter configured

	show ip filter configured Displays IPv4 filtering configuration.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000

Example	switch (config) # show ip filter configured
	
	Packet filtering for IPv4: enabled
	
	IPv4 configuration:
	  Chain 'input' Policy 'accept':
	  Rule 1:
	  Target : accept
	  Protocol : all
	  Source : all
	  Destination : 1.1.1.0/24
	  Interface : all
	  State : any
	  Other Filter: -
	
	  Chain 'output' Policy 'accept':
	  Rule 1:
	  Target : reject-with icmp-net-unreachable
	  Protocol : tcp
	  Source : all
	  Destination : all
	  Interface : all
	  State : any
	  Other Filter: dest-port 1000

Related Commands	
Notes	

4.1.9.2.9 show ipv6 filter

	show ipv6 filter Displays IPv6 filtering state.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000

Example	switch (config) # show ipv6 filter
	
	Packet filtering for IPv6: enables
	
	Active IPv6 filtering rules (omitting any not from configuration):
	  Chain 'input' Policy 'accept':
	  Rule 1:
	  Target : accept
	  Protocol : all
	  Source : all
	  Destination : 1.1.1.0/24
	  Interface : all
	  State : any
	  Other Filter: -
	
	  Chain 'output' Policy 'accept':
	  Rule 1:
	  Target : reject-with icmp-net-unreachable
	  Protocol : tcp
	  Source : all
	  Destination : all
	  Interface : all
	  State : any
	  Other Filter: dest-port 1000

Related Commands	
Notes	

4.1.9.2.10 show ipv6 filter all

	show ipv6 filter all Displays IPv6 filtering state (including un-configured rules).
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000

Example	switch (config) # show ipv6 filter all
	
	Packet filtering for IPv6: enables
	
	All active IPv6 filtering rules:
	  Chain 'input' Policy 'accept':
	  Rule 1:
	  Target : accept
	  Protocol : all
	  Source : all
	  Destination : 1.1.1.0/24
	  Interface : all
	  State : any
	  Other Filter: -
	
	  Chain 'output' Policy 'accept':
	  Rule 1:
	  Target : reject-with icmp-net-unreachable
	  Protocol : tcp
	  Source : all
	  Destination : all
	  Interface : all
	  State : any
	  Other Filter: dest-port 1000

Related Commands	
Notes	

4.1.9.2.11 show ipv6 filter configured

	show ipv6 filter configured Displays IPv6 filtering configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000

Example	`switch (config) # show ipv6 filter configured`
	
	`Packet filtering for IPv6: enables`
	
	`IPv6 configuration:`
	`  Chain 'input' Policy 'accept':`
	`  Rule 1:`
	`  Target  : accept`
	`  Protocol  : all`
	`  Source  : all`
	`  Destination : 1.1.1.0/24`
	`  Interface  : all`
	`  State  : any`
	`  Other Filter:  -`
	
	`  Chain 'output' Policy 'accept':`
	`  Rule 1:`
	`  Target  : reject-with icmp-net-unreachable`
	`  Protocol  : tcp`
	`  Source  : all`
	`  Destination : all`
	`  Interface  : all`
	`  State  : any`
	`  Other Filter: dest-port 1000`

Related Commands	
Notes	

4.2 Chassis Management

The chassis manager provides the user access to the following information:

Accessible Parameters	Description
switch temperatures	Displays system's temperature
power supply voltages	Displays power supplies' voltage levels
fan unit	Displays system fans' status
power unit	Displays system power consumers
Flash memory	Displays information about system memory utilization.

Additionally, it monitors:

- AC power to the PSUs
- DC power out from the PSUs
- Chassis failures

4.2.1 System Health Monitor

The system health monitor scans the system to decide whether or not the system is healthy. When the monitor discovers that one of the system's modules (fan, or power supply) is in an unhealthy state or returned from an unhealthy state, it notifies the users through the following methods:

- System logs—accessible to the user at any time as they are saved permanently on the system
- Status LEDs—changed by the system health monitor when an error is found in the system and is resolved
- Email/SNMP traps—notification on any error found in the system and resolved

4.2.1.1 Re-Notification on Errors

When the system is in an unhealthy state, the system health monitor notifies the user about the current unresolved issue every X seconds. The user can configure the re-notification gap by running the “health notif-cntr <counter>” command.

4.2.1.2 System Health Monitor Alerts Scenarios

System Health Monitor sends notification alerts in the following cases:

Alert Message	Scenario	Notification Indicator	Recovery Action	Recovery Message
<fan_name> speed is below minimal range	A chassis fan speed is below minimal threshold (15% of maximum speed)	Email, fan LED and system status LED set red, log alert, SNMP.	Check the fan and replace it if required	“<fan_name> has been restored to its normal state”
<fan_name> is unresponsive	A chassis fan is not responsive on the switch system	Email, fan LED and system status LED set red, log alert, SNMP	Check fan connectivity and replace it if required	“<fan_name> has been restored to its normal state”
<fan_name> is not present	A chassis fan is missing	Email, fan LED and system status LED set red, log alert, SNMP	Insert a fan unit	“<fan_name> has been restored to its normal state”
Insufficient number of working fans in the system	Insufficient number of working fans in the system	Email, fan LED and system status LED set red, log alert, SNMP	Plug in additional fans or change faulty fans	“The system currently has sufficient number of working fans”
Power Supply <ps_number> voltage is out of range	The power supply voltage is out of range.	Email, power supply LED and system status LED set red, log alert, SNMP	Check the power connection of the PS	“Power Supply <ps_number> voltage is in range”
Power supply <ps_number> temperature is too hot	A power supply unit temperature is higher than the maximum threshold of 70 Celsius on the switch system	Email, power supply LED and system status LED set red, log alert, SNMP	Check chassis fans connections. On switch systems, check system fan connections.	“Power supply <ps_number> temperature is back to normal”

Alert Message	Scenario	Notification Indicator	Recovery Action	Recovery Message
Power Supply <number> is unresponsive	A power supply is malfunctioning or disconnected	Email, system status and power supply LED set red, log alert, SNMP	Connect power cable or replace malfunctioning PS	“Power supply has been removed” or “PS has been restored to its normal state”
ASIC temperature is too hot	An ASIC unit temperature is higher than the maximum threshold of 105 Celsius on switch systems	Email, system status LED set red, log alert, SNMP	Check the fan’s system	“ASIC temperature is back to normal”

4.2.2 Power Management

4.2.2.1 Width Reduction Power Saving

Link width reduction (LWR) is a
NVIDIA

proprietary power saving feature to be utilized to economize the power usage of the fabric. LWR may be used to manually or automatically configure a certain connection between
NVIDIA switch

systems to lower the width of a link from 4X operation to 1X based on the traffic flow.

LWR is relevant only for
40GbE

speeds in which the links are operational at a 4X width.

> When “show interfaces” is used, a port’s speed appears unchanged even when only one lane is active.

LWR has three operating modes per interface:

- Disabled—LWR does not operate and the link remains in 4X under all circumstances.
- Automatic—the link automatically alternates between 4X and 1X based on traffic flow.
- Force—a port is forced to operate in 1X mode lowering the throughput capability of the port. This mode should be chosen in cases where constant low throughput is expected on the port for a certain time period—after which the port should be configured to one of the other two modes, to allow higher throughput to pass through the port.

The following table describes LWR configuration behavior:

Switch-A Configuration	Switch-B Configuration	Behavior
Disable	Disable	LWR is disabled

Switch-A Configuration	Switch-B Configuration	Behavior
Disable	Force	Transmission from Switch-B to Switch-A operates at 1X. On the opposite direction, LWR is disabled.
Disable	Auto	Depending on traffic flow, transmission from Switch-B to Switch-A may operate at 1X. On the opposite direction, LWR is disabled.
Auto	Force	Transmission from Switch-B to Switch-A operates at 1 lane. Transmission from Switch-A to Switch-B may operate at 1X depending on the traffic.
Auto	Auto	Width of the connection depends on the traffic flow
Force	Force	Connection between the switches operates at 1x

4.2.3 Monitoring Environmental Conditions

1. Display module's temperature. Run:

```
switch (config) # show temperature
-------------------------------------------------------
Module      Component              Reg  CurTemp    Status
                                        (Celsius)
-------------------------------------------------------
MGMT        SIB                    T1   33.00      OK
MGMT        Board AMB temp         T1   24.50      OK
MGMT        Ports AMB temp         T1   27.00      OK
MGMT        CPU package Sensor     T1   29.00      OK
MGMT        CPU Core Sensor        T1   28.00      OK
MGMT        CPU Core Sensor        T2   24.00      OK
PS1         power-mon              T1   22.00      OK
PS2         power-mon              T1   23.00      OK
```

2. Display measured voltage levels of power supplies. Run:

```
switch (config) # show voltage
-------------------------------------------------------------------------------------------------------
Module   Power Meter             Reg                    Expected  Actual    Status  High   Low
                                                        Voltage   Voltage           Range  Range
-------------------------------------------------------------------------------------------------------
MGMT     acdc-monitor1           DDR3 0.675V            0.68      0.67      OK      0.78   0.57
MGMT     acdc-monitor1           CPU 0.9V               0.78      0.78      OK      0.89   0.66
MGMT     acdc-monitor1           SYS 3.3V               3.30      3.34      OK      3.79   2.80
MGMT     acdc-monitor1           CPU 1.8V               1.80      1.79      OK      2.07   1.53
MGMT     acdc-monitor1           CPU/PCH 1.05V          1.05      1.05      OK      1.21   0.89
MGMT     acdc-monitor1           CPU 1.05V              1.05      1.05      OK      1.21   0.89
MGMT     acdc-monitor1           DDR3 1.35V             1.35      1.35      OK      1.55   1.15
MGMT     acdc-monitor1           USB 5V                 5.00      5.04      OK      5.75   4.25
MGMT     acdc-monitor1           1.05V LAN              1.50      1.50      OK      1.72   1.27
MGMT     ASICVoltMonitor1        Asic 1.2V              1.20      1.21      OK      1.38   1.02
MGMT     ASICVoltMonitor1        Asic 3.3V              3.30      3.32      OK      3.79   2.80
MGMT     ASICVoltMonitor2        Vcore SPC              0.95      0.96      OK      1.09   0.81
MGMT     acdc-monitor2           1.8V Switch SPC        1.80      1.82      OK      2.07   1.53
PS1      power-mon               N/A                    0.00      0.00      FAIL    0.00   0.00
PS2      power-mon               vout 12V               12.00     11.98     OK      13.80  10.20
```

3. Display the fan speed and status. Run:

```
switch (config) # show fan
-----------------------------------------------------
Module          Device          Fan  Speed      Status
                                     (RPM)
-----------------------------------------------------
FAN1            FAN             F1   9305.00    OK
FAN2            FAN             F1   8823.00    OK
FAN3            FAN             F1   9057.00    OK
FAN4            FAN             F1   9369.00    OK
PS1             FAN             F1   10288.00   OK
PS2             FAN             -    -          NOT PRESENT
```

4. Display the voltage current and status of each module in the system. Run:

```
switch (config) # show power consumers
-------------------------------------------------------------------
Module  Device            Sensor  Power   Voltage  Current  Status
                                  [Watts] [Volts]  [Amp]
-------------------------------------------------------------------
PS1     power-mon         input   37.50   12.02    3.19     OK
MGMT    acdc-monitor2     input   -       -        -        OK

Total power used : 37.50 Watts
```

4.2.4 USB Access

The OS can access USB devices attached to switch systems. USB devices are automatically recognized and mounted upon insertion. To access a USB device for reading or writing a file, you need to provide the path to the file on the mounted USB device in the following format:

```
scp://username:password@hostname/var/mnt/usb1/<file name>
```

While username and password are the admin username and password and hostname is the IP of the switch.

Examples:

- To fetch an image from a USB device, run the command:

```
switch (config) # image fetch scp://username:password@hostname/var/mnt/usb1/<image filename>
```

- To save log file (my-logfile) to a USB device under the name “test_logfile” using the command “logging files”, run:

```
switch (config) # logging files upload my-logfile scp://username:password@hostname/var/mnt/usb1/
test_logfile
```

- To safely remove the USB and to flush the cache, after writing (log files, for example) to a USB, use the “usb eject” command:

```
switch (config) # usb eject
```

4.2.5 Unit Identification LED

The unit identification (UID) LED is a hardware feature used as a means of locating a specific switch system in a server room.

To activate the UID LED on a switch system, run:

```
switch (config) # led MGMT uid on
```

To verify the LED status, run:

```
switch (config) # show leds
Module   LED            Status
------------------------------------------------------------------------------
MGMT     UID            Blue
```

To deactivate the UID LED on a switch system, run:

```
switch (config) # led MGMT uid off
```

4.2.6 System Reboot

To reboot your switch system, run:

```
switch (config) # reload
```

4.2.7 Viewing Active Events

NVIDIA Onyx supports viewing all active events on the system. The following events may be observed with the command "show system hardware events".

Event Name	Description
Ethernet Family	
Invalid Mac (SMAC=MC)	Source MAC is a multicast address
Invalid Mac (SMAC=DMAC)	Source MAC is same as destination mac address
Invalid Ethertype	Packet has an unknown Ethertype (0x05DC < ethertype < 0x600)
IP Routing Family	
Ingress Router interface is disabled	Ingress packet has been dropped because incoming L3 interface is admin down
Mismatched IP (UC DIP over MC/BC Mac)	Packet MAC is multicast/broadcast but destination IP is unicast
Invalid IP (DIP=loopback)	Destination IP is loopback IP (For IPv6: DIP==::1/128 or DIP==0:0:0:0:0:ffff:7f00:0/104 For IPv4: DIP==127.0.0.0/8)
Invalid IP (SIP=MC)	Source IP is multicast address (For IPv6: SIP == FF00::/8 For IPv4: SIP == 224.0.0.0: 239.255.255.255 aka 224.0.0.0/4)
Invalid IP (SIP=unspecified)	Source IP is unspecified
Invalid IP (SIP=DIP)	Source IP is identical to destination IP
Mismatched MC Mac	Packet's multicast MAC does not correspond to packet's MC IP address
IPv6 neighbor not resolved	IPv6 neighbor not resolved
Invalid IPv6 (SIP=Link Local)	Source IP is link local (IPv6)
MC RPF check failure	Multicast RPF check failure
TTL expired	TTL value is zero
Egress Router interface is disabled	Egress packet has been dropped because outgoing L3 interface is admin/oper is down

IPv4 neighbor not resolved	Entry not found for destination
Tunnel Family	
NVE Decap fragmentation error	Fragmentation error during decapsulation

4.2.8 Chassis Management Commands

- 4.2.8.1 Chassis Management
  - 4.2.8.1.1 clear counters
  - 4.2.8.1.2 clear system hardware events
  - 4.2.8.1.3 health
  - 4.2.8.1.4 led uid
  - 4.2.8.1.5 power enable
  - 4.2.8.1.6 system profile
  - 4.2.8.1.7 usb eject
  - 4.2.8.1.8 show asic-version
  - 4.2.8.1.9 show bios
  - 4.2.8.1.10 show cpld
  - 4.2.8.1.11 show fan
  - 4.2.8.1.12 show health-report
  - 4.2.8.1.13 show inventory
  - 4.2.8.1.14 show leds
  - 4.2.8.1.15 show memory
  - 4.2.8.1.16 show module
  - 4.2.8.1.17 show power
  - 4.2.8.1.18 show power consumers
  - 4.2.8.1.19 show protocols
  - 4.2.8.1.20 show resources
  - 4.2.8.1.21 show system capabilities
  - 4.2.8.1.22 show system hardware events
  - 4.2.8.1.23 show system mac
  - 4.2.8.1.24 show system profile
  - 4.2.8.1.25 show system profile detailed
  - 4.2.8.1.26 show system type
  - 4.2.8.1.27 show temperature
  - 4.2.8.1.28 show version
  - 4.2.8.1.29 show version concise
  - 4.2.8.1.30 show voltage

4.2.8.1 Chassis Management

4.2.8.1.1 clear counters

	clear counters [all | interface <type> <number>] Clears switch counters.	
Syntax Description	all	Clears all switch counters.
	type	A specific interface type.
	number	The interface number.
Default	N/A	
Configuration Mode	config	
History	3.2.3000 3.6.4000: Added note	
Example	`switch (config) # clear counters`	
Related Commands		
Notes	The command also clears storm-control counters	

4.2.8.1.2 clear system hardware events

	clear system hardware events Clears all active events.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000
Example	`switch (config) # clear system hardware events`
Related Commands	show system hardware events
Notes	

4.2.8.1.3 health

	health {max-report-len <length> | re-notif-cntr <counter> | report-clear} Configures health daemon settings.	
Syntax Description	max-report-len <length>	Sets the length of the health report (number of line entries) Range: 10-2048
	re-notif-cntr <counter>	Health control changes notification counter in seconds Range: 120-7200
	report-clear	Clears the health report
Default	max-report-len: 50 re-notif-cntr:	
Configuration Mode	config	

History	3.1.0000
Example	`switch (config) # health re-notif-cntr 125`
Related Commands	show health-report
Notes	

4.2.8.1.4 led uid

	led <module> uid <on | off> Configures the UID LED.	
Syntax Description	module	Specifies the module whose UID LED to configure
	on	Turns on UID LED
	off	Turns off UID LED
Default	N/A	
Configuration Mode	config	
History	3.6.1002 3.6.2002:	
Example	`switch (config) # led MGMT uid on`	
Related Commands		
Notes		

4.2.8.1.5 power enable

	power enable <module name> no power enable <module name> Powers on the module. The no form of the command shuts down the module.	
Syntax Description	module name	Enables power for selected module
Default	Power is enabled on all modules	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # power enable L01`	
Related Commands	show power show power consumers	
Notes	• It is recommended to run this command prior to extracting a module from the switch system, else errors are printed in the log	

4.2.8.1.6 system profile

	system profile {eth-default | eth-ipv6-max | eth-ipv4-mc-max} [force] Optimizes switch system profile to preferred mode.	
Syntax Description	eth-default	Balanced Ethernet profile
	eth-ipv6-max	Optimized profile for IPv6 scale
	eth-ipv4-mc-max	Optimized profile for IPv4 multicast scale
	force	Forces operation, without the need for user confirmation
Default	eth-default	
Configuration Mode	config	
History	3.6.6000	
Example	`switch (config) # system profile eth-default`	
Related Commands	show system profile	
Notes		

4.2.8.1.7 usb eject

	usb eject Turns off the USB interface gracefully.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Example	`switch (config) # usb eject`
Related Commands	
Notes	Applicable only for systems with USB interface.

4.2.8.1.8 show asic-version

	show asic-version Displays firmware ASIC version.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.4.2008: Updated example

Example	`switch (config) # show asic-version`
	`================================================`
	`Module  Device  Version`
	`================================================`
	`MGMT  SPC  15.0200.0092`

Related Commands	
Notes	

4.2.8.1.9 show bios

	show bios Displays the BIOS version information.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.4150

Example	`switch (config) # show bios`
	`BIOS version : 4.6.5`
	`BIOS subversion : Official AMI Release`
	`BIOS release date : 07/02/2021`

Related Commands	
Notes	

4.2.8.1.10 show cpld

	show cpld Displays status of all CPLDs in the system.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.3.4302: Updated example 3.10.1000: Updated example to reflect the part number (PN) field 3.10.1100: Updated example to reflect Version Minor

Example	`switch (config) # show cpld`
	`-----------------------------------------------------------------------------`
	`Name            Type                 Version          Version Minor   PN`
	`-----------------------------------------------------------------------------`
	`Cpld1           CPLD_TOR             9                1               0x0078`
	`Cpld2           CPLD_SWB_UNIFIED     3                3               0x0128`
	`Cpld3           CPLD_LED             1                0               0x00d1`

Related Commands	
Notes	

4.2.8.1.11 show fan

	show fan Displays fans status.
Syntax Description	N/A
Default	N/A

Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config) # show fan --- Module Device Fan Speed Status (RPM) --- FAN1 FAN F1 9305.00 OK FAN2 FAN F1 8823.00 OK FAN3 FAN F1 9057.00 OK FAN4 FAN F1 9369.00 OK PS1 FAN F1 10288.00 OK PS2 FAN - - NOT PRESENT
Related Commands	
Notes	

4.2.8.1.12 show health-report

	show health-report Displays health report.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.3.0000: Output update
Example	switch (config) # show health-report ======================== | ALERTS CONFIGURATION | ======================== Re-notification counter (sec):[3600] Report max counter: [50] ======================== | HEALTH REPORT | ======================== No Health issues file
Related Commands	health
Notes	

4.2.8.1.13 show inventory

	show inventory Displays system inventory.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.4.1604: Removed CPU module output from example 3.5.1000: Removed Type column from example 3.6.1002: Updated example
Example	

```
switch (config) # show inventory
---------------------------------------------------------------------
Module      Part Number         Serial Number        Asic Rev.    HW Rev.
---------------------------------------------------------------------
CHASSIS     MSN2100-CB2F        MT1752X06330         N/A          B3
MGMT        MSN2100-CB2F        MT1752X06330         1            B3
```

Related Commands	
Notes	

4.2.8.1.14 show leds

	show leds [<module>] Displays the LED status of the switch system.	
Syntax Description	module	Specifies the module whose LED status to display
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002 3.6.2002: Updated example	
Example	switch (config) # show leds Module LED Status --- MGMT STATUS Green MGMT FAN1 Green MGMT FAN2 Green MGMT FAN3 Green MGMT FAN4 Green MGMT PS_STATUS Green MGMT PS1 Green MGMT PS2 Green MGMT UID Blue	
Related Commands		
Notes		

4.2.8.1.15 show memory

	show memory Displays memory status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.7.1000: Updated example
Example	

```
switch (config) # show memory

------------------------------------------------------------------------
Memory Space  Total       Used        Free        Used+B/C    Free-B/C
------------------------------------------------------------------------
Physical      15848 MB    2849 MB     12999 MB    3854 MB     11994 MB
Swap          0 MB        0 MB        0 MB

Physical Memory Borrowed for System Buffers and Cache:
  Buffers            : 27 MB
  Cache              : 910 MB
  Total Buffers/Cache: 937 MB
```

Related Commands	
Notes	

4.2.8.1.16 show module

	show module Displays modules status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.3.0000: Added “Is Fatal” column 3.4.2008: Updated command output 3.4.3000: Updated command output and added note

Example	`switch (config) # show module`
	`======================`
	` Module    Status`
	`======================`
	` MGMT      ready`
	` FAN1      ready`
	` FAN2      ready`
	` PS1       ready`
	` PS2       not-present`

Related Commands	
Notes	The Status column may have one of the following values: error, fatal, not-present, powered-off, powered-on, ready.

4.2.8.1.17 show power

	show power Displays power supplies and power usage.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.5.1000: Updated example
Example	

```
switch (config) # show power
-------------------------------------------------------------------------------
Module   Device      Sensor Power   Voltage   Current   Capacity   Feed   Status
                         [Watts] [Volts]   [Amp]      [Watts]
-------------------------------------------------------------------------------
PS1      power-mon  input  32.25   12.11     1.26       800.00     DC     OK
PS2      power-mon  input  46.56   12.13     2.33       800.00     DC     OK
```

Related Commands	
Notes	

4.2.8.1.18 show power consumers

	show power consumers Displays power consumption information.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.5.1000: Updated example

Example

```
switch (config) # show power consumers
-------------------------------------------------------------------------
Module  Device            Sensor  Power   Voltage  Current  Status
                                  [Watts] [Volts]  [Amp]
-------------------------------------------------------------------------
MGMT    CURR_MONITOR      12V     52.96   11.71    4.52     OK
PS1     power-mon         input   252.00  12.00    20.25    OK
PS2     power-mon         input   280.00  12.03    23.25    OK

Total power used : 52.96 Watts
```

Related Commands	
Notes	

4.2.8.1.19 show protocols

	show protocols Displays all protocols enabled in the system.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.2.3000 3.3.4550: Updated example 3.6.1002: Updated example

Example	switch (config) # show protocols Ethernet enabled spanning-tree rst lacp disabled lldp disabled igmp-snooping disabled ets enabled priority-flow-control disabled sflow disabled openflow disabled mlag disabled dot1x disabled isolation-group disabled IP routing disabled bgp disabled pim disabled vrrp disabled ospf disabled magp disabled dhcp-relay disabled
Related Commands	
Notes	

4.2.8.1.20 show resources

	show resources Displays system resources.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config) # show resources Total Used Free Physical 2027 MB 761 MB 1266 MB Swap 0 MB 0 MB 0 MB Number of CPUs: 1 CPU load averages: 0.11 / 0.23 / 0.23 CPU 1 Utilization: 5% Peak Utilization Last Hour: 19% at 2012/02/15 13:26:19 Avg. Utilization Last Hour: 7%
Related Commands	
Notes	

4.2.8.1.21 show system capabilities

	show system capabilities Displays system capabilities.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode

History	3.1.0000 3.3.0000: Added gateway support 3.6.1002: Updated example 3.7.0000: Updated example

Example	`switch (config) # show system capabilities`
	`Ethernet: Supported, L2, L3`
	`Ethernet Max licensed speed: 100Gb`

Related Commands	show system profile
Notes	

4.2.8.1.22 show system hardware events

	show system hardware events <family-name> [clear-on-read] Displays all active events.	
Syntax Description	family-name	Displays all active events per event family: • ethernet • tunnel • ip
	clear-on-read	Clears all active events after displaying them
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6000	

Example	`switch (config) # show system hardware events clear-on-read`
	
	`Ethernet:  smac is mc;`
	`  smac equal dmac;`
	`IP:  packet to router is not ip;`
	`Tunnel:`	

Related Commands		
Notes		

4.2.8.1.23 show system mac

	show system mac Displays system MAC address.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	`switch (config) # show system mac` `00:02:c9:5e:AF:18`
Related Commands	N/A
Notes	

4.2.8.1.24 show system profile

	show system profile Displays system profile.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.2.0000 3.7.0000: Updated example

Example	`switch (config) # show system profile`
	
	`Profile: eth-default`

Related Commands	system profile
Notes	

4.2.8.1.25 show system profile detailed

	show system profile detailed Displays detailed system profile.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.6000

Example	`switch (config) # show system profile detailed`
	
	`Profile: eth-default`
	
	`----------------------------------------------`
	`Parameter Guaranteed Max Value`
	`----------------------------------------------`
	`FDB size 102400`
	`IPMC-L2 lists 10240`
	`IPMC-L3 lists 10240`
	`IPv4 MC/IGMP routes 10240`
	`IPv4 neighbors 51200`
	`IPv6 neighbors 8192`
	`IPv4 routes 100000`
	`IPv6 shorts 51200`
	`IPv6 routes 21504`
	`VRF 64`
	`RIF 999`

Related Commands	system profile
Notes	

4.2.8.1.26 show system type

	show system type Displays system type.
Syntax Description	N/A

Default	N/A
Configuration Mode	Any command mode
History	3.5.1000
Example	switch (config) # show system type SN2100
Related Commands	
Notes	

4.2.8.1.27 show temperature

	show temperature Displays system temperature sensors status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config) # show temperature ---------- Module Component Reg CurTemp Status (Celsius) ---------- MGMT SPC T1 43.00 OK MGMT Ports AMB temp T1 31.00 OK MGMT Board AMB temp T1 30.00 OK MGMT CPU Core Sensor T1 23.00 OK MGMT CPU Core Sensor T2 23.00 OK MGMT CPU Core Sensor T3 24.00 OK MGMT CPU Core Sensor T4 24.00 OK
Related Commands	
Notes	

4.2.8.1.28 show version

	show version Displays version information for the currently running system image.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000

Example	switch (config) # show version Product name: Onyx Product release: 3.6.8008 Build ID: #1-dev Build date: 2018-07-18 13:46:44 Target arch: x86_64 Target hw: x86_64 Built by: jenkins@c5de6027485e Version summary: X86_64 3.6.8008 2018-07-18 13:46:44 x86_64 Product model: x86 Host ID: 7CFE9058E01E System UUID: 03000200-0400-0500-0006-000700080009 Uptime: 16h 50m 41.260s CPU load averages: 2.38 / 2.25 / 2.24 Number of CPUs: 2 System memory: 2860 MB used / 12988 MB free / 15848 MB total Swap: 0 MB used / 0 MB free / 0 MB total
Related Commands	
Notes	

4.2.8.1.29 show version concise

	show version concise Displays concise version information for the currently running system image.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config) # show version concise X86_64 3.6.4006 2017-07-03 16:17:39 x86_64
Related Commands	
Notes	

4.2.8.1.30 show voltage

	show voltage Displays voltage level measurements on different sensors.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.3.5006: Updated example
Example	

```
switch (config) # show voltage

================================================================================================
Module  Power Meter           Reg                  Expected  Actual   Status  High   Low
                                                   Voltage   Voltage          Range  Range
================================================================================================
MGMT    BOARD_MONITOR         USB 5V sensor        5.00      5.15     OK      5.55   4.45
MGMT    BOARD_MONITOR         Asic I/O sensor      2.27      2.11     OK      2.55   1.99
MGMT    BOARD_MONITOR         1.8V sensor          1.80      1.79     OK      2.03   1.57
MGMT    BOARD_MONITOR         SYS 3.3V sensor      3.30      3.28     OK      3.68   2.92
MGMT    BOARD_MONITOR         CPU 0.9V sensor      0.90      0.93     OK      1.04   0.76
MGMT    BOARD_MONITOR         1.2V sensor          1.20      1.19     OK      1.37   1.03
MGMT    CPU_BOARD_MONITOR     12V sensor           12.00     11.67    OK      13.25  10.75
MGMT    CPU_BOARD_MONITOR     12V sensor           2.50      2.46     OK      2.80   2.20
MGMT    CPU_BOARD_MONITOR     2.5V sensor          3.30      3.26     OK      3.68   2.92
MGMT    CPU_BOARD_MONITOR     SYS 3.3V sensor      3.30      3.24     OK      3.68   2.92
MGMT    CPU_BOARD_MONITOR     SYS 3.3V sensor      1.80      1.79     OK      2.03   1.57
MGMT    CPU_BOARD_MONITOR     1.8V sensor          1.20      1.24     OK      1.37   1.03
```

Related Commands	
Notes	

4.3 Management Source IP Address

4.3.1 ntp source-interface

In many cases network operators prefer to have a single IP address for the switch that is used for management operations like switch configuration, receiving remote log files, ping, and so forth. That IP address is needed for building firewall rules so that network switches can be easily identified. It is also required for identifying management traffic and exact management target in network logs.

The following protocols are supported by the feature:

- FTP
- TFTP
- NTP
- Syslog
- TACACS
- SSH, SSHD, SCP
- Ping
- Traceroute
- SNMP

4.3.2 Commands

4.3.2.1 ssh server listen

	ssh server listen <interface> no ssh server listen <interface> Defines a source interface for ssh server.	
Syntax Description	interface	Interface to bind. Possible values: mgmt0, lo, or loopback 0-31

Default	N/A
Configuration Mode	config
History	3.7.1002
Example	`switch (config)# ssh server listen loopback2`
Related Commands	
Notes	

4.3.2.2 ssh client global source-interface

	ssh client global [vrf <vrf-name>] source-interface <interface> no ssh client global [vrf <vrf-name>] source-interface <interface> Configures the source interface that binds the SSH client to a specific address used by the slogin command.	
Syntax Description	interface	Interface to bind. Possible values: loopback0-31
	vrf-name	VRF to be affected. If "vrf-name" parameter is not specified, "default" VRF will be used.
Default	N/A	
Configuration Mode	config	
History	3.7.1002	
	3.9.2000	Added VRF option
Example	`switch (config)# ssh client global source-interface loopback10`	
Related Commands		
Notes	The <interface> must be in the <vrf-name>. Source-interface could be configured in any VRF that the configured service is enabled in.	

4.3.2.3 ip ftp source-interface

	ip ftp [vrf <vrf-name>] source-interface <interface> no ip ftp [vrf <vrf-name>] source-interface <interface> Configures the source interface for ftp protocol. The no form of the command disables the ftp source interface protocol.	
Syntax Description	interface	Interface to bind Possible values: loopback0-31
	vrf-name	VRF to be affected. If "vrf-name" parameter is not specified, "default" VRF will be used.
Default	N/A	
Configuration Mode	config	
History	3.7.1002	
	3.9.2000	Added VRF option
Example	`switch (config)# ip ftp source-interface loopback7`	
Related Commands		

Notes	The <interface> must be in the <vrf-name>. The source-interface can be configured for each existing VRF.

4.3.2.4 ip tftp source-interface

	ip tftp [vrf <vrf-name>] source-interface <interface> no ip tftp [vrf <vrf-name>] source-interface <interface> Configures the source interface for tftp protocol. The no form of the command disables the tftp source interface protocol.	
Syntax Description	interface	Interface to bind. Possible values: loopback0-31
	vrf-name	VRF to be affected. If "vrf-name" parameter is not specified, "default" VRF will be used.
Default	N/A	
Configuration Mode	config	
History	3.7.1002	
	3.9.2000	Added VRF option
Example	`switch (config)# ip tftp source-interface loopback7`	
Related Commands		
Notes	The <interface> must be in the <vrf-name>. The source-interface can be configured for each existing VRF.	

4.3.2.5 ip scp source-interface

	ip scp [vrf <vrf-name>] source-interface <interface> no ip scp [vrf <vrf-name>] source-interface <interface> Configures the source interface for scp protocol. The no form of the command disables the scp source interface protocol.	
Syntax Description	interface	Interface to bind Possible values: loopback0-31
	vrf-name	VRF to be affected. If "vrf-name" parameter is not specified, "default" VRF will be used.
Default	N/A	
Configuration Mode	config	
History	3.8.1000	
	3.9.2000	Added VRF option
Example	`switch (config)# ip scp source-interface loopback7`	
Related Commands		
Notes	The <interface> must be in the <vrf-name>. The source-interface can be configured for each existing VRF.	

4.3.2.6 ip sftp source-interface

	ip sftp [vrf <vrf-name>] source-interface <interface> no ip sftp [vrf <vrf-name>] source-interface <interface> Configures the source interface for sftp protocol. The no form of the command disables the sftp source interface protocol.	
Syntax Description	interface	Interface to bind Possible values: loopback0-31
	vrf-name	VRF to be affected. If "vrf-name" parameter is not specified, "default" VRF will be used.
Default	N/A	
Configuration Mode	config	
History	3.8.1000	
	3.9.2000	Added VRF option
Example	`switch (config)# ip sftp source-interface loopback7`	
Related Commands		
Notes	The <interface> must be in the <vrf-name>. The source-interface can be configured for each existing VRF.	

4.3.2.7 ip traceroute source-interface

	ip traceroute source-interface <interface> no ip traceroute source-interface <interface> Configures the source interface for traceroute protocol. The no form of the command disables the traceroute source interface protocol.	
Syntax Description	interface	Interface to bind Possible values: loopback0-31
Default	N/A	
Configuration Mode	config	
History	3.8.1000	
	3.9.2000	Added VRF option
Example	`switch (config)# ip traceroute source-interface loopback7`	
Related Commands		
Notes	The <interface> must be in the <vrf-name>. The source-interface can be configured for each existing VRF.	

4.3.2.8 logging source-interface

	logging [vrf <vrf-name>] soucrce-interface <interface> no logging [vrf <vrf-name>] source-interface <interface> Configures the source interface for sending the log messages to remote servers. The no form of the command disables the logging source interface protocol.

Syntax Description	interface	Interface to bind Possible values: loopback0-31
	vrf-name	VRF to be affected. If "vrf-name" parameter is not specified, "default" VRF will be used.
Default	N/A	
Configuration Mode	config	
History	3.7.1002	
	3.9.0600	Added notes
	3.9.2000	Added VRF option
Example	`switch (config)# logging source-interface loopback7`	
Related Commands		
Notes	• Source interface is supported only for logging host using UDP and not supported for TCP • Changes in runtime in the dns regarding a logging host (changes of relation between hostname and ip) are not handled, logging source ip may stop working • The <interface> must be in the <vrf-name>. Source-interface could be configured in any VRF that the configured service is enabled in.	

4.3.2.9 tacacs-server source-interface

	tacacs-server [vrf <vrf-name>] source-interface <interface> no tacacs-server [vrf <vrf-name>] source-interface <interface> Configures the source interface for tacacs protocol. The no form of the command disables the tacacs source interface protocol.	
Syntax Description	interface	Interface to bind Possible values: loopback0-31
	vrf-name	VRF to be affected. If "vrf-name" parameter is not specified, "default" VRF will be used.
Default	N/A	
Configuration Mode	config	
History	3.7.1002	
	3.9.2000	Added VRF option
Example	`switch (config)# tacacs source-interface loopback23`	
Related Commands		
Notes	The <interface> must be in the <vrf-name>. Source-interface must be in the same VRF that the configured service is enabled in.	

4.3.2.10 ip icmp source-interface

	ip icmp [vrf <vrf-name>] source-interface <interface> no ip icmp [vrf <vrf-name>] source-interface <interface> Configures the source interface for icmp protocol (for ping requests). The no form of the command disables the icmp source interface protocol.

Syntax Description	interface	Interface to bind Possible values: loopback0-31
	vrf-name	VRF to be affected. If "vrf-name" parameter is not specified, "default" VRF will be used.
Default	N/A	
Configuration Mode	config	
History	3.7.1002	
	3.9.2000	Added VRF option
Example	`switch (config)# ip icmp source-interface loopback24`	
Related Commands		
Notes	<interface> must be in the <vrf-name>. Source-interface can be configured for each existing VRF.	

4.3.2.11 ntp source-interface

	ntp [vrf <vrf-name>] soucrce-interface <interface> no ntp [vrf <vrf-name>] source-interface <interface> Configures the source interface for ntp protocol. This interface will be used for user requested and periodic ntp synchronization. The no form of the command disables the ntp source interface protocol.	
Syntax Description	interface	Interface to bind. Range: loopback0-31.
	vrf-name	VRF to be affected. If "vrf-name" parameter is not specified, "default" VRF will be used.
Default	N/A	
Configuration Mode	config	
History	3.7.1002	
	3.9.2000	Added VRF option
Example	`switch (config)# ntp source-interface loopback7`	
Related Commands		
Notes	• This command sets source IP for NTPD and NTP date • The <interface> must be in the <vrf-name>. Source-interface must be in the same VRF that the configured service is enabled in.	

4.3.2.12 snmp-server source-interface

	snmp-server [vrf <vrf-name>] source-interface <interface> no snmp-server [vrf <vrf-name>] source-interface <interface> Configures the source interface for sending SNMP traps and informs. The no form of the command disables the snmp-server source interface protocol.	
Syntax Description	interface	Interface to bind Range: loopback0-31
	vrf-name	VRF to be affected. If "vrf-name" parameter is not specified, "default" VRF will be used.
Default	N/A	

Configuration Mode	config	
History	3.8.1000	
	3.9.2000	Added VRF option
Example	`switch (config) # snmp-server source-interface loopback7`	
Related Commands	show snmp source-interface	
Notes	The <interface> must be in the <vrf-name>. Source-interface could be configured in any VRF that the configured service is enabled in.	

4.3.2.13 show ip ftp source-interface

	show ip ftp [vrf {<vrf-name>|all}] source-interface Displays the source interface.	
Syntax Description	vrf-name	Describes VRF that will be affected by this command. If "vrf" parameter is not specified, the "default" VRF will be used implicitly.
Default	N/A	
Configuration Mode	Any configuration mode	
History	3.7.1002	
	3.9.2000	Added VRF option

Example	`switch (config)# show ip ftp source-interface`
	`Source IP for ftp client:`
	`Configured: loopback7`
	`Current : loopback7`
	`IPv4-addr : 5.5.5.5`
	`IPv6-addr : none`
	
	`switch (config)# show ip ftp vrf all source-interface`
	
	`VRF name: default`
	
	`Source IP for ftp client:`
	`  Configured: none`
	`  Current  : none`
	`  IPv4-addr : none`
	`  IPv6-addr : none`
	
	`VRF name: mgmt`
	
	`Source IP for ftp client:`
	`  Configured: loopback2`
	`  Current  : loopback2`
	`  IPv4-addr : 10.10.10.10`
	`  IPv6-addr : none`	

Related Commands		
Notes		

4.3.2.14 show ip tftp source-interface

	show ip tftp [vrf {<vrf-name>|all}] source-interface Displays the source interface.

Syntax Description	vrf-name	Describes VRF that will be affected by this command. If "vrf" parameter is not specified, the "default" VRF will be used implicitly.
Default	N/A	
Configuration Mode	Any configuration mode	
History	3.9.2000	

Example	switch (config)# show ip tftp [vrf {<vrf-name>|all}] source-interface
	
	Example:
	show ip tftp vrf all source-interface
	
	VRF name: default
	
	Source IP for tftp client:
	  Configured: none
	  Current : none
	  IPv4-addr : none
	  IPv6-addr : none
	
	VRF name: mgmt
	
	Source IP for tftp client:
	  Configured: loopback2
	  Current : loopback2
	  IPv4-addr : 10.10.10.10
	  IPv6-addr : none	

Related Commands		
Notes		

4.3.2.15 show ntp source-interface

	show ntp [vrf {<vrf-name>|all}] source-interface Displays the source interface.	
Syntax Description	vrf-name	Describes VRF that will be affected by this command. If "vrf" parameter is not specified, the "default" VRF will be used implicitly.
Default	N/A	
Configuration Mode	Any configuration mode	
History	3.7.1002	
	3.9.2000	Added VRF option

Example	switch (config)# show ntp source-interface
	Source IP for ntp client:
	Configured: loopback2
	Current : loopback2
	IPv4-addr : 10.7.144.97
	IPv6-addr : none
	
	switch (config)# show ntp vrf all source-interface
	
	VRF name: default
	
	Source IP for ntp client:
	  Configured: none
	  Current : none
	  IPv4-addr : none
	  IPv6-addr : none
	
	VRF name: mgmt
	
	Source IP for ntp client:
	  Configured: loopback2
	  Current : loopback2
	  IPv4-addr : 10.10.10.10
	  IPv6-addr : none

Related Commands	
Notes	

4.3.2.16 show logging source-interface

	show logging [vrf {<vrf-name>|all}] source-interface Displays the source interface.	
Syntax Description	vrf-name	Describes VRF that will be affected by this command. If "vrf" parameter is not specified, the "default" VRF will be used implicitly.
Default	N/A	
Configuration Mode	Any configuration mode	
History	3.7.1002	
	3.9.2000	Added VRF option

Example	switch (config)# show logging source-interface
	Source IP for syslogd client:
	Configured: loopback23
	Current : loopback23
	IPv4-addr : 1.3.5.7
	IPv6-addr : none
	
	switch (config)# show logging vrf all source-interface
	
	VRF name: default
	
	Source IP for syslogd client:
	  Configured: none
	  Current : none
	  IPv4-addr : none
	  IPv6-addr : none
	
	VRF name: mgmt
	
	Source IP for syslogd client:
	  Configured: loopback2
	  Current : loopback2
	  IPv4-addr : 10.10.10.10
	  IPv6-addr : none	

Related Commands		

Notes	

4.3.2.17 show tacacs source-interface

	show tacacs [vrf {<vrf-name>|all}] source-interface Displays the source interface.	
Syntax Description	vrf-name	Describes VRF that will be affected by this command. If "vrf" parameter is not specified, the "default" VRF will be used implicitly.
Default	N/A	
Configuration Mode	Any configuration mode	
History	3.7.1002	
	3.9.2000	Added VRF option.

Example	`switch (config)# show tacacs source-interface`
	`Source IP for tacacs client:`
	`Configured: loopback3`
	`Current : loopback3`
	`IPv4-addr : 1.3.5.7`
	`IPv6-addr : none`
	
	`switch (config)# show tacacs vrf all source-interface`
	
	`VRF name: default`
	
	`Source IP for tacacs client:`
	`  Configured: none`
	`  Current  : none`
	`  IPv4-addr : none`
	
	`VRF name: mgmt`
	
	`Source IP for tacacs client:`
	`  Configured: loopback2`
	`  Current  : loopback2`
	`  IPv4-addr : 10.10.10.10`	

Related Commands		
Notes		

4.3.2.18 show ip icmp source-interface

	show ip icmp [vrf {<vrf-name>|all}] source-interface Displays the source interface.	
Syntax Description	vrf-name	Describes VRF that will be affected by this command. If "vrf" parameter is not specified, the "default" VRF will be used implicitly.
Default	N/A	
Configuration Mode	Any configuration mode	
History	3.7.1002	
	3.9.2000	Added VRF option

Example	switch (config)# show icmp source-interface Source IP for ping client: Configured: none Current : none IPv4-addr : none IPv6-addr : none switch (config)# show ip icmp vrf all source-interface VRF name: default Source IP for ping client: Configured: none Current : none IPv4-addr : none IPv6-addr : none VRF name: mgmt Source IP for ping client: Configured: loopback2 Current : loopback2 IPv4-addr : 10.10.10.10 IPv6-addr : none
Related Commands	
Notes	

4.3.2.19 show ip traceroute source-interface

	show ip traceroute [vrf {<vrf-name>|all}] source-interface Displays the source interface.	
Syntax Description	vrf-name	VRF that will be affected by this command. If "vrf" parameter is not specified, the "default" VRF will be used implicitly.
Default	N/A	
Configuration Mode	Any configuration mode	
History	3.7.1002	
	3.9.2000	Added VRF option
Example	switch (config)# show traceroute source-interface Source IP for traceroute client: Configured: none Current : none IPv4-addr : none IPv6-addr : none switch (config)# show ip traceroute vrf all source-interface VRF name: default Source IP for traceroute client: Configured: none Current : none IPv4-addr : none IPv6-addr : none VRF name: mgmt Source IP for traceroute client: Configured: loopback2 Current : loopback2 IPv4-addr : 10.10.10.10 IPv6-addr : none	
Related Commands		

Notes	

4.3.2.20 show ssh client source-interface

	show ssh client [vrf {<vrf-name>|all}] source-interface Displays the SSH client source interface.	
Syntax Description	vrf-name	Describes VRF that will be affected by this command. If "vrf" parameter is not specified, the "default" VRF will be used implicitly.
Default	N/A	
Configuration Mode	Any configuration mode	
History	3.7.1002	
	3.7.1100	Updated example
	3.9.2000	Added VRF option

Example	`switch (config)# show ssh client source-interface`
	`Source IP for ssh client:`
	`Configured: loopback1`
	`Current : loopback1`
	`IPv4-addr : 1.1.1.1`
	`IPv6-addr : none`
	
	`switch (config)# show ssh client vrf all source-interface`
	
	`VRF name: default`
	
	`Source IP for ssh client:`
	`Configured: none`
	`Current : none`
	`IPv4-addr : none`
	`IPv6-addr : none`
	
	`VRF name: mgmt`
	
	`Source IP for ssh client:`
	`Configured: loopback2`
	`Current : loopback2`
	`IPv4-addr : 10.10.10.10`
	`IPv6-addr : none`	

Related Commands		
Notes		

4.3.2.21 show ip scp source-interface

	show ip scp [vrf {<vrf-name>|all}] source-interface Displays the source interface.	
Syntax Description	vrf-name	Describes VRF that will be affected by this command. If "vrf" parameter is not specified, the "default" VRF will be used implicitly.
Default	N/A	
Configuration Mode	Any configuration mode	
History	3.7.1002	
	3.9.2000	Added VRF option.

Example	switch (config)# show ip scp source-interface Source IP for scp client: Configured: none Current : none IPv4-addr : none IPv6-addr : none switch (config)# show ip scp vrf all source-interface VRF name: default Source IP for scp client: Configured: none Current : none IPv4-addr : none IPv6-addr : none VRF name: mgmt Source IP for scp client: Configured: loopback2 Current : loopback2 IPv4-addr : 10.10.10.10 IPv6-addr : none
Related Commands	
Notes	

4.3.2.22 show ip sftp source-interface

	show ip sftp [vrf {<vrf-name>|all}] source-interface Displays the source interface.	
Syntax Description	vrf-name	Describes VRF that will be affected by this command. If "vrf" parameter is not specified, the "default" VRF will be used implicitly.
Default	N/A	
Configuration Mode	Any configuration mode	
History	3.7.1002	
	3.9.2000	Added VRF option
Example	switch (config)# show ip sftp source-interface Source IP for sftp client: Configured: none Current : none IPv4-addr : none IPv6-addr : none switch (config)# show ip sftp vrf all source-interface VRF name: default Source IP for sftp client: Configured: none Current : none IPv4-addr : none IPv6-addr : none VRF name: mgmt Source IP for sftp client: Configured: loopback2 Current : loopback2 IPv4-addr : 10.10.10.10 IPv6-addr : none	

Related Commands	
Notes	

4.3.2.23 show snmp source-interface

	show snmp [vrf {<vrf-name>|all}] source-interface Displays the source interface for sending SNMP traps and informs.	
Syntax Description	vrf-name	Describes VRF that will be affected by this command. If "vrf" parameter is not specified, the "default" VRF will be used implicitly.
Default	N/A	
Configuration Mode	config	
History	3.8.1000	
	3.9.2000	Added VRF option

Example	`switch (config)# show snmp source-interface`
	`Source IP for snmp server:`
	
	`Configured: loopback7`
	`Current : loopback7`
	`IPv4-addr : 5.5.5.5`
	`IPv6-addr : none`
	
	`switch (config)# show snmp vrf all source-interface`
	
	`VRF name: default`
	
	`Source IP for snmp server:`
	`  Configured: none`
	`  Current  : none`
	`  IPv4-addr : none`
	`  IPv6-addr : none`
	
	`VRF name: mgmt`
	
	`Source IP for snmp server:`
	`  Configured: loopback2`
	`  Current  : loopback2`
	`  IPv4-addr : 10.10.10.10`
	`  IPv6-addr : none`	

Related Commands	snmp-server source-interface <interface>	
Notes		

4.4 Upgrade/Downgrade Process

The following pages provide information on upgrading and downgrading the operating system version on the device.

- Important Pre-OS Upgrade Notes
- Upgrading Operating System Software
- Upgrading HA Groups
- Upgrading MLAG-STP Setup
- Deleting Unused Images

- Downgrading OS Software
- Upgrading System Firmware
- Image Maintenance Using ONIE
- Software Management Commands

4.4.1 Important Pre-OS Upgrade Notes

Please consider the following items prior to upgrading the operating system:

- The system becomes unavailable while OS upgrade is in progress
- The upgrade procedure burns the software image as well as the firmware should there be a need
- Before upgrading the software image on your system, make sure to close all CLI sessions besides the one used to run the upgrade process
- To upgrade the NVIDIA Onyx on an MLAG cluster, please refer to "Upgrading HA Groups"
- When upgrading from a version older than 3.6.3130 with an MLAG cluster, "show mlag" output appears as "UP" and "Peering" state instead of "Upgrade" on both MLAG VIP clusters—the upgrade process will not be affected
- Interfaces with global pause are not mapped to a lossless pool after upgrade from versions earlier than 3.6.5000
- You have to read and accept the End-User License Agreement (EULA) after image upgrade in case the EULA is modified. The EULA link is only available upon first login to CLI
- Linux docker container names are limited to 180 characters. Upgrading to this version removes containers which do not comply with this limitation and prints the following warning to the log: "Removed configuration of container: <container name>, container name is limited to 180 characters"

4.4.2 Upgrading Operating System Software

To upgrade NVIDIA Onyx, perform the following steps.

1. Enter Config mode.

```
switch > enable
switch # configure terminal
switch (config) #
```

2. Display the currently available image (.img file).

```
switch (config) # show images
Installed images:

  Partition 1:
  <old_image>

  Partition 2:
  <old_image>

Last boot partition: 1
Next boot partition: 1

Images available to be installed:
 webimage.tbz
  <old_image>

Serve image files via HTTP/HTTPS: no

No image install currently in progress.

Boot manager password is set.
```

```
Image signing: trusted signature always required
Admin require signed images: yes

Settings for next boot only:
   Fallback reboot on configuration failure: yes (default)
```

3. Delete the image listed under “Images available to be installed” prior to fetching the new image. Use the command “image delete” for this purpose.

```
switch (config) # image delete <old_image>
```

> When deleting an image, it is recommended to delete the file, but not the partition, so as to not overload system resources.

4. Fetch the new software image.

```
switch (config) # image fetch scp://<username>:<password>@<ip-address>/var/www/html/<new_image>
Password (if required): ****** 100.0%[########################################################################]
```

5. Display the available images again and verify that the new image now appears under “Images available to be installed”.

> To recover from image corruption (e.g., due to power interruption), there are two installed images on the system. See the commands “image boot next” and “image boot location” for more information.

```
switch (config) # show images
Installed images:

  Partition 1:
  <old_image>

  Partition 2:
  <old_image>

Last boot partition: 1
Next boot partition: 1

Images available to be installed:
 webimage.tbz
  <new_image>

Serve image files via HTTP/HTTPS: no

No image install currently in progress.

Boot manager password is set.

Image signing: trusted signature always required
Admin require signed images: yes

Settings for next boot only:
   Fallback reboot on configuration failure: yes (default)
```

6. Install the new image.

```
switch (config) # image install <new_image>
Step 1 of 4: Verify Image
 100.0%  [############################################################]
Step 2 of 4: Uncompress Image
 100.0%  [############################################################]
Step 3 of 4: Create Filesystems
 100.0%  [############################################################]
Step 4 of 4: Extract Image
 100.0%  [############################################################]
```

CPU utilization may go up to 100% during image upgrade.

7. Have the new image activate during the next boot.

```
switch (config) # image boot next
```

8. Run “show images” to review your images.

```
switch (config) # show images
Installed images:

  Partition 1:
  <new_image>

  Partition 2:
  <old_image>

Last boot partition: 1
Next boot partition: 1

Images available to be installed:
 webimage.tbz
  <new_image>

Serve image files via HTTP/HTTPS: no

No image install currently in progress.

Boot manager password is set.

Image signing: trusted signature always required
Admin require signed images: yes

Settings for next boot only:
   Fallback reboot on configuration failure: yes (default)
```

9. Save current configuration.

```
switch (config) # configuration write
```

10. Reboot to run the new image.

```
switch (config) # reload
Configuration has been modified; save first? [yes] yes
Configuration changes saved.
Rebooting...
switch (config)#
```

After software reboot, the software upgrade will also automatically upgrade the firmware version.

When performing an upgrade from the WebUI, make sure that the image being upgraded to is not already located in the system (i.e., fetched from the CLI).

4.4.3 Upgrading HA Groups

If fallback is ever necessary in an HA group, all cluster nodes must have the same OS version installed and they must be immediately reloaded.

To upgrade NVIDIA Onyx version without affecting an HA group:

1. Identify the HA group master.
   For MLAG. Run:

```
switch (config)# show mlag-vip
MLAG VIP
========
MLAG group name: my-mlag-group
MLAG VIP address: 10.234.23.254 /24
Active nodes: 2

Hostname             VIP-State            IP Address
-----------------------------------------------------
SwitchA              master               10.234.23.1
SwitchB              standby              10.234.23.254
```

2. Upgrade standby node in the HA group according to steps 1-10 in "Upgrading Operating System Software".
3. Wait until all standby nodes have rejoined the group.

   In situations of heavy CPU load or noisy network, it is possible that another node assumes the role of cluster master before all standby nodes have rejoined the group. If this happens, you may stop waiting and proceed directly to step 4.

   When slave upgrade is complete and the master is still in the lower version, MACs are not learned by the slave switch system (except for traffic flood) until master switch upgrade is complete.

4. Upgrade the master node in the HA group according to steps 1-10 in "Upgrading Operating System Software".

4.4.4 Upgrading MLAG-STP Setup

To upgrade NVIDIA Onyx on an MLAG-STP setup from 3.6.610x to this version, there are two possible procedures:

Procedure 1

1. Make sure there are no loops in the fabric.
2. Disable STP. Run:

```
switch (config) # no spanning-tree
```

3. Perform the upgrade according to steps 1-10 in "Upgrading Operating System Software".
4. Enable STP – this step may lead to traffic loss while the STP state is converging. Run:

```
switch (config) # spanning-tree
```

Procedure 2:

1. Shutdown all ports except for the MLAG IPL Port-Channel on the MLAG stanby switch, make sure not to shutdown the MLAG IPL Port-Channel ports.
2. Save configuration. Run:

```
switch (config) # configuration write
```

3. Upgrade MLAG stanby switch according to steps 1-10 in "Upgrading Operating System Software".

4. Once the MLAG stanby switch is back online with the new version, use "show mlag" and "show mlag-vip" commands and verify both processes are up. Next enable all ports on the MLAG stanby switch, and verify all ports are back online.
5. Upgrade MLAG master according to steps 1-3 above, and use the below command to reboot the master:

```
switch (config) # reload force immediate
```

6. Once the MLAG master is back online with the new version, use "show mlag" and "show mlag-vip" command and verify all is up. Next enable all ports on the MLAG master, and verify all ports are back online.

4.4.5 Deleting Unused Images

To delete unused images, conduct the following steps.

1. Get a list of the unused images.

```
switch (config) # show images

Installed images:
  Partition 1:
    version: image-X86_64-3.6.5000.img

  Partition 2:
    version: image-X86_64-3.6.5000.img

Last boot partition: 1
Next boot partition: 1

Images available to be installed:
  No image files are available to be installed.

Serve image files via HTTP/HTTPS: no

No image install currently in progress.
Boot manager password is set.

Image signing              : trusted signature always required
Admin require signed images: yes

Settings for next boot only:
  Fallback reboot on configuration failure: yes (default)
```

2. Delete the unused images.

```
switch (config) # image delete image-X86_64-3.9.1302.img
```

> When deleting an image, it is recommended to delete the file, but not the partition, so as to not overload system resources.

4.4.6 Downgrading OS Software

Prior to downgrading software, please make sure the following prerequisites are met.

1. Log in to the switch via the CLI using the console port.
2. Backup configuration by following these steps.

a. Disable paging of CLI output.

```
switch (config)# no cli default paging enable
```

b. Display commands to recreate current running configuration.

```
switch (config)# show running-config
```

c. Copy the output to a text file.

4.4.6.1 Downloading Image

1. Log in to your system to obtain its product number.

```
switch (config) # show inventory
```

2. Log in to NVIDIA Enterprise Support Portal and download the relevant NVIDIA Onyx version to your system type
3. Log in to your system via the CLI.
4. Change to Config mode.

```
switch > enable
switch # configure terminal
switch (config) #
```

5. Delete all previous images from the Images available to be installed prior to fetching the new image.
6. Fetch the desired software image.

```
switch (config) # image fetch scp://username:password@192.168.10.125/var/www/html/<image_name>
100.0%[######################################################################]
```

4.4.6.2 Downgrading Image

> The procedure described below assumes that booting and running is done from Partition 1 and the downgrade procedure is performed on Partition 2.

1. Log in to your system via the CLI as admin.
2. Enter config mode.

```
switch > enable
switch # configure terminal
```

3. Display all image files on the system.

```
switch (config) # show images
Images available to be installed:
new_image.img
 <downgrade version> 2010-09-19 16:52:50
Installed images:
```

```
Partition 1:
 <current version> 2010-09-19 03:46:25
Partition 2:
 <current version> 2010-09-19 03:46:25
Last boot partition: 1
Next boot partition: 1
No boot manager password is set.
```

4. Install the fetched image.

```
switch (config) # image install <image_name>
Step 1 of 4: Verify Image
100% [#################################################################]
Step 2 of 4: Uncompress Image
100.0% [################################################################]
Step 3 of 4: Create Filesystems
100.0% [################################################################]
Step 4 of 4: Extract Image
100.0% [################################################################]
```

5. Display all image files on the system.

```
switch (config) # show images
Images available to be installed:
new_image.img
 <downgrade version> 2010-09-19 16:52:50
Installed images:
Partition 1:
 <current version> 2010-09-19 03:46:25
Partition 2:
 <downgrade version> 2010-09-19 16:52:50
Last boot partition: 1
Next boot partition: 2
No boot manager password is set.
```

6. Configure the boot location to be the other (next) partition.

```
switch (config) # image boot next
```

There are two installed images on the system. Therefore, if one of the images gets corrupted (due to power interruption, for example), in the next reboot the image will go up from the second partition.

If you are downgrading to an older software version which has never been run yet on the switch, use the following command sequence as well.

```
switch (config) # no boot next fallback-reboot enable
switch (config) # configuration write
```

7. Reload.

```
switch (config) # reload
```

4.4.6.3 Switching to Partition with Older Software Version

The system saves a backup configuration file when upgrading from an older software version to a newer one. If the system returns to the older software partition, it uses this backup configuration file.

> All configuration changes done with the new software are lost when returning to the older software version.

There are 2 instances where the backup configuration file does not exist:

- The user has run “reset factory” command, which clears all configuration files in the system
- The user has run “configuration switch-to” to a configuration file with different name then the backup file

> Note that the configuration file becomes empty if the system is downgraded to a software version which has never been installed yet.

To allow switching partition to the older software version for the 2 aforementioned cases only, follow the steps below.

1. Run the following command.

```
switch (config)# no boot next fallback-reboot enable
```

2. Set the boot partition.

```
switch (config)# image boot next
```

3. Save the configuration.

```
switch (config)# configuration write
```

4. Reload the system.

```
switch (config)# reload
```

4.4.7 Upgrading System Firmware

NVIDIA Onyx software package version has a default switch firmware version. When you update the operating system software to a new version, an automatic firmware update process will be attempted by NVIDIA Onyx. This process is described below.

4.4.7.1 After Updating Software

Upon rebooting your switch system after updating the OS software, the OS compares its default firmware version with the currently programmed firmware versions on all the switch modules.
If one or more of the switch modules is programmed with a firmware version other than the default version, then the OS automatically attempts to burn the default firmware version instead.

> If a firmware update takes place, then the login process is delayed a few minutes.

To verify that the firmware update was successful, log into your switch and run the command “show asic-version” (can be run in any mode). This command lists all of the switch modules along with their firmware versions. Make sure that all the firmware versions are the same and match the default firmware version. If the firmware update failed for one or more modules, then the following warning is displayed.
Some subsystems are not updated with a default firmware.

> If you detect a mismatch in firmware version for one or more modules of the switch system, please contact your assigned field application engineer.

4.4.7.2 Importing Firmware and Changing the Default Firmware

To perform an automatic firmware update by the OS for a different switch firmware version without changing the OS version, import the firmware package as described below. The OS sets it as the new default firmware and performs the firmware update automatically as described in the previous subsections.

4.4.7.2.1 Default Firmware Change on Standalone Systems

1. Import the firmware image (.mfa file). Run:

```
switch (config) # image fetch scp://root@1.1.1.1:/tmp/fw-SIB-rel-11_1600_0200-FIT.mfa
Password (if required): *******
100.0% [##############################################################################]
switch (config) # image default-chip-fw fw-SIB-rel-11_1600_0200-FIT.mfa
Installing default firmware image. Please wait...
Default Firmware 11.1600.0200 updated. Please save configuration and reboot for new FW to take effect.
```

2. Save the configuration. Run:

```
switch (config) # configuration write
```

3. Reboot the system to enable auto update.

4.4.8 Image Maintenance Using ONIE

ONIE is an “open compute” Open Network Install Environment for bare metal network switches. ONIE enables a bare metal network switch ecosystem where end-users have a choice among different network operating systems.

NVIDIA Onyx is distributed in way that allows installation on an ONIE environment. Certain switch models come pre-installed with ONIE and NVIDIA Onyx and support changing to a different operating system (OS).

To change the switch system's operating OS:

1. Reboot the switch and wait for it to reach the GRUB menu:

```
                                GNU GRUB version 2.02

X86_64 3.4.1932 2015-04-24 18:04:12 x86_64 1
X86_64 3.4.1932 2015-04-24 18:04:12 x86_64 2
ONIE
```

2. Select the ONIE option using the arrow keys. The following message appears:

```
Due to security constraints, this option will uninstall your current MLNX OS system.
Are you sure ?
```

3. Type YES to continue.
   Since NVIDIA Onyx is being uninstalled and deleted from the hard drive, the process takes a few hours. After this is finished, the system reboots into the ONIE shell and auto discovery begins.

```
Info: Fetching tftp://<ip-address>/7C-FE-90-5E-6A-4A/onie-installer-x86_64-mlnx_x86-r5.0.1400 ...
Failure: Unable to find installer: /installer
Info: Fetching tftp://<ip-address>/0AE016FB/onie-installer-x86_64-mlnx_x86-r5.0.1400 ...
Failure: Unable to find installer: /installer
Info: Fetching tftp://<ip-address>/0AE016F/onie-installer-x86_64-mlnx_x86-r5.0.1400 ...
...
```

4. In order to manually insert an install URL, press Enter and insert the command "onie-nos-install <http> / <tftp> <url> <image name .bin>". For example:

```
onie-nos-install http://<ip_address>//sx_mlnx_os-3.5.1000-21/X86_64/X86_64-3.5.1000-21-installer.bin
```

   Once you hit Enter, you have about 4 second to insert the command so it is recommended to prepare the command in advance and simply pasting it in. At this stage, the OS installation begins.
5. Wait for the installation to end and reboot this switch to boot into the OS.

```
ONIE:/ # onie-nos-install http://<ip_address>//sx_mlnx_os-3.5.1000-21/X86_
64/X86_64-3.5.1000-21-installer.bin
Stopping: discover... done.
down.
ONIE: eth1: link down.  Skipping configuration.
ONIE: Failed to configure eth1 interface
Info: Fetching http://<ip_address>//sx_mlnx_os-3.5.1000-21/X86_64/X86_64-3.5.1000-21-installer.bin ...
Connecting to <ip_address>
installer           100% |*******************************|   392M  0:00:00 ETA
ONIE: Executing installer: http://<ip_address>//sx_mlnx_os-3.5.1000-21/X86_64/X86_64-3.5.1000-21-
installer.bin
```

4.4.9 Software Management Commands

- 4.4.9.1 image boot
- 4.4.9.2 boot next
- 4.4.9.3 boot system
- 4.4.9.4 image default-chip-fw
- 4.4.9.5 image delete

- 4.4.9.6 image fetch
- 4.4.9.7 image install
- 4.4.9.8 image move
- 4.4.9.9 image options
- 4.4.9.10 show bootvar
- 4.4.9.11 show images

4.4.9.1 image boot

	image boot {location <location-ID> | next} Specifies the default location where the system should be booted from.	
Syntax Description	location-ID	Specifies the default destination location. There can be up to 2 images on the system. The possible values are 1 or 2.
	next	Sets the boot location to be the next once after the one currently booted from, thus avoiding a cycle through all the available locations.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # image boot location 2`	
Related Commands	show images	
Notes		

4.4.9.2 boot next

	boot next fallback-reboot enable no boot next fallback-reboot enable Sets the default setting for next boot. Normally, if the system fails to apply the configuration on startup (after attempting upgrades or downgrades, as appropriate), it will reboot to the other partition as a fallback. The no form of the command tells the system not to do that, only for the next boot.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.2.0506
Example	`switch (config) # boot next fallback-reboot enable`
Related Commands	show images

Notes	• Normally, if the system fails to apply the configuration on startup (after attempting upgrades or downgrades, as appropriate) it reboots to the other partition as a fallback. • The no form of this command tells the system not to do that only for the next boot. In other words, this setting is not persistent and goes back to being enabled automatically after each boot. • When downgrading to an older software version which has never been run yet on a system, the “fallback reboot” always happens, unless the command “no boot next fallback-reboot enable” is used. However, this also happens when the older software version has been run before, but the configuration file has been switched since upgrading. In general, a downgrade only works (without having the fallback reboot forcibly disabled) if the process can find a snapshot of the configuration file (by the same name as the currently active one) which was taken before upgrading from the older software version. If that is not found, a fallback reboot is performed in preference to falling back to the initial database because the latter generally involves a loss of network connectivity, and avoiding that is of paramount importance.

4.4.9.3 boot system

	boot system {location | next} no boot system next Configures which system image to boot by default. The no form of the command resets the next boot location to the current active one.	
Syntax Description	location	Specifies location from which to boot system • 1—installs to location 1 • 2—installs to location 2
	next	Boots system from next location after one currently booted
Default	N/A	
Configuration Mode	config	
History	3.2.0506	
Example	`switch (config) # boot system location 2`	
Related Commands	show images	
Notes		

4.4.9.4 image default-chip-fw

	image default-chip-fw <filename> no image default-chip-fw <original-fw-filename> Sets the default firmware package to be installed. The no form of the command resets default firmware package.	
Syntax Description	filename	Specifies the firmware filename
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.6.6000	Added the no form of the command

Example	`switch (config) # image default-chip-fw <filename>.mfa`
Related Commands	show asic-version show images
Notes	

4.4.9.5 image delete

	image delete <image-name> Deletes the specified image file.	
Syntax Description	image-name	Specifies the image name
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # image delete <filename>.img`	
Related Commands	show images	
Notes		

4.4.9.6 image fetch

	`image fetch [vrf <vrf-name>] <URL> [<filename>]` Downloads an image from the specified URL or via SCP.	
Syntax Description	vrf-name—Describes docker daemon VRF context, impacts fetching images and running containers. If "vrf" parameter is not specified, the "default" VRF will be used.	
	URL	HTTP, HTTPS, FTP, TFTP, SCP and SFTP are supported Example: scp://username[:password]@hostname/path/filename
	filename	Specifies a filename for this image to be stored as locally
Default	N/A	
Configuration Mode	config	
History	3.1.0000 3.9.2000—Added VRF option	
Example		

```
switch (config) # image fetch scp://<username>@192.168.10.125/var/www/html/<image_name>
Password ******
100.0%[############################################################]
switch (config) #
```

Other options:

```
switch (config) # image fetch http://10.1.0.40/path/filename
switch (config) # image fetch http://[fd4f:13:cc00:1::40]/path/filename
switch (config) # image fetch ftp://user:mypassword@10.1.0.40/foo/bar.img
switch (config) # image fetch ftp://user:mypassword@[fd4f:13:cc00:1::40]/foo/bar.img
switch (config) # image fetch tftp://hostname/dir/filename
switch (config) # image fetch tftp://[fd4f:13:cc00:1::40]/dir/filename
switch (config) # image fetch scp://user@myhost/dir/filename
switch (config) # image fetch scp://user@myhost:1022/dir/filename
switch (config) # image fetch scp://user:pass@[fd4f:13:cc00:1::40]/dir/filename
switch (config) # image fetch sftp://user@myhost/dir/filename
switch (config) # image fetch sftp://user@[fd4f:13:cc00:1::40]:1022/dir/filename
switch (config) # image fetch sftp://user:pass@[fd4f:13:cc00:1::40]/dir/filename
```

Related Commands	show images
Notes	• Please delete the previously available image, prior to fetching the new image • The path to the file in the case of TFTP depends on the server configuration. Therefore, it may not be an absolute path but a relative one. • See "Upgrading Operating System Software" page

4.4.9.7 image install

	image install <image-filename> [location <location-ID>] | [progress <prog-options>] Installs the specified image file.	
Syntax Description	image-filename	Specifies the image name
	location-ID	Specifies the image destination location
	prog-options	• “no-track” overrides CLI default and does not track the installation progress • “track” overrides CLI default and tracks the installation progress
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # image install X86_64 3.6.5000 2017-07-26 06:54:12 x86_64 Step 1 of 4: Verify Image 100.0% [##] Step 2 of 4: Uncompress Image 100.0% [##] Step 3 of 4: Create Filesystems 100.0% [##] Step 4 of 4: Extract Image 100.0% [##] switch (config) #	
Related Commands	show images	
Notes	• The image cannot be installed on the “active” location (the one which is currently being booted) • On a two-location system, the location is chosen automatically if no location is specified	

4.4.9.8 image move

	image move <src-image-name> <dest-image-name> Renames the specified image file.	
Syntax Description	src-image-name	Specifies the current image name
	dest-image-name	Specifies the new image name
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # image move image1.img image2.img`	
Related Commands	show images	
Notes		

4.4.9.9 image options

	image options serve all no image options serve all Configures options and defaults for image usage. The no form of the command disables options and defaults for image usage.	
Syntax Description	serve all	Specifies that the image files present on this appliance should be made available for HTTP and/or HTTPS download
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # image options serve all`	
Related Commands	show images	
Notes	The parameter “serve all” affects not only the files currently present, but also any files that are later downloaded. It only applies to image files, not the installed images, which are not themselves in a downloadable format. After running “serve all” the URLs where the images will be available are: • http://<HOSTNAME>/system_images/<FILENAME> • https://<HOSTNAME>/system_images/<FILENAME>	

4.4.9.10 show bootvar

	show bootvar Displays the installed system images and the boot parameters.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000

Example	switch (config)# show bootvar Installed images: Partition 1: X86_64 3.6.4110-12 2017-07-26 06:54:12 x86_64 Partition 2: X86_64 3.6.4006 2017-07-03 16:17:39 x86_64 Last boot partition: 1 Next boot partition: 1 Serve image files via HTTP/HTTPS: no Boot manager password is set. Image signing: trusted signature always required Admin require signed images: yes Settings for next boot only: Fallback reboot on configuration failure: yes (default)
Related Commands	
Notes	

4.4.9.11 show images

	show images Displays information about the system images and boot parameters.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config)# show images Installed images: Partition 1: X86_64 3.6.4110-12 2017-07-26 06:54:12 x86_64 Partition 2: X86_64 3.6.4006 2017-07-03 16:17:39 x86_64 Last boot partition: 1 Next boot partition: 1 Images available to be installed: webimage.tbz X86_64 3.6.4071-12 2017-07-26 06:54:12 x86_64 Serve image files via HTTP/HTTPS: no No image install currently in progress. Boot manager password is set. Image signing: trusted signature always required Admin require signed images: yes Settings for next boot only: Fallback reboot on configuration failure: yes (default)
Related Commands	show images
Notes	

4.5 Configuration Management

4.5.1 Saving a Configuration File

To save the current configuration to the active configuration file, you can either use the “configuration write” command (requires running in Config mode) or the “write memory” command (requires running in Enable mode).

- To save the configuration to the active configuration file, run:

```
switch (config) # configuration write
```

- To save the configuration to a user-specified file without making the new file the active configuration file, run:

```
switch (config) # configuration write to myconf no-switch
```

- To save the configuration to a user-specified file and make the new file the active configuration file, run:

```
switch (config) # configuration write to myconf
```

- To display the available configuration files and the active file, run:

```
switch (config) # show configuration files
initial
myconf (active)
switch (config) #
```

4.5.2 Loading a Configuration File

By default, or after a system reset, the system loads the default “initial” configuration file.

To load a different configuration file and make it the active configuration:

```
switch >
switch > enable
switch # configure terminal
switch (config) # configuration switch-to myconfig
switch (config) #
```

4.5.3 Restoring Factory Default Configuration

If system configuration becomes corrupted, it is suggested to restore factory default configuration.

- To restore factory default configuration on a single management module system, run:

```
switch (config) # reset factory keep-basic
```

4.5.4 Managing Configuration Files

There are two types of configuration files that can be applied on the BIN files (binary) and text-based configuration files.

4.5.4.1 BIN Configuration Files

BIN configuration files are not human readable. Additionally, these files are encrypted and contain integrity verification preventing them from being edited and used.

- To create a new BIN configuration file, do the following:

```
switch (config) # configuration new my-filename
```

> A newly created BIN configuration file is always empty and is not created from the running-config.

- To upload a BIN configuration file to an external file server, do the following:

```
switch (config) # configuration upload my-filename scp://myusername@my-server/path/to/my/<file>
```

- To fetch a BIN configuration file, do the following:

```
switch (config) # configuration fetch scp://myusername@my-server/path/to/my/<file>
```

- To see the available configuration files, do the following:

```
switch (config) # show configuration files
initial (active)
my-filename

Active configuration: initial
Unsaved changes:      no
switch (config) #
```

- To load a BIN configuration file, do the following:

```
switch (config) # configuration switch-to my-filename
This requires a reboot.
Type 'yes' to confirm: yes
```

> A binary configuration file uploaded from the switch is encrypted and has integrity verification. If the file is modified in any manner, the fetch to the switch fails.

4.5.4.2 Text Configuration Files

Text configuration files are text-based and editable. It is similar in form to the output of the command “show running-config expanded”.

- To create a new text-based configuration file, do the following:

```
switch (config) # configuration text generate active running save my-filename
```

A newly created text configuration file is always created from the running-config.

- To apply a text-based configuration file, do the following:

```
switch (config) # configuration text file my-filename apply
```

```
switch (config) # configuration text generate active running save my-filename
```

Applying a text-based configuration file to an existing/running data port configuration may result in unpredictable behavior. It is therefore suggested to first clear the configuration by applying a specific configuration file (following the procedure in "BIN Configuration File") or by resetting the switch back to factory default.

- To upload a text-based configuration file to an external file server, do the following:

```
switch (config) # configuration text file my-filename upload scp://root@my-server/root/tmp/my-filename
```

- To fetch a text-based configuration file from an external file server to a switch, do the following:

```
switch (config) # configuration text fetch scp://root@my-server/root/tmp/my-filename
```

- To apply a text-based configuration file, do the following:

```
switch (config) # configuration text file my-filename apply
```

When applying a text-based configuration file, the configuration is appended to the existing configuration. Only new or changed configuration is added. Reboot is not required.

4.5.5 Automated Periodic Configuration File Backup

4.5.5.1 Automated Backup

Automated configuration file backup feature can be used to upload the active configuration file on every “configuration write".

- To set the remote URL to upload the configuration file to, run the following:

```
switch (config) # configuration auto-upload remote-url “scp://root:password@my-server/path/to/upload/to”
```

- To check the remote URL set, run the following:

```
switch (config) # show configuration auto-upload
    Auto-upload settings:
        Enabled:      yes
        Remote url:   scp://root@my-server/path/to/upload/to
        Password  :   ******
```

- To save the configuration, run the following:

```
switch (config)# configuration write
```

This will upload the active configuration file on every “configuration write."

- To remove the remote URL, run the following:

```
switch (config)# no configuration auto-upload remote-url
```

This will disable the feature. It will not upload the active configuration file after each “configuration write."

4.5.5.2 Automated Periodic Backup

Scheduled jobs can be used to perform automated periodic backup.

To upload the active configuration file periodically, follow these steps.

1. Create a job.

```
switch (config) # job 1
```

2. Add the upload command to the job.

```
switch (config) # job 1 command 1 "configuration upload timestamp active scp://root:password@my-server/
path/to/upload/to"
```

3. Schedule this job to run periodically, and specify the period.

```
switch (config) # job 1 schedule periodic interval 18h0m0s
```

4. Enable the job.

```
switch (config) # job 1 enable
```

4.5.6 Configuration Management Commands

- 4.5.6.1 File System
    - 4.5.6.1.1 debug generate dump
    - 4.5.6.1.2 file debug-dump
    - 4.5.6.1.3 file stats
    - 4.5.6.1.4 file tcpdump
    - 4.5.6.1.5 file eula upload
    - 4.5.6.1.6 file open-source-licenses upload

- 4.5.6.1.7 file help-docs upload
- 4.5.6.1.8 reload
- 4.5.6.1.9 reset factory
- 4.5.6.1.10 configuration new factory
- 4.5.6.1.11 configuration new factory keep-docker
- 4.5.6.1.12 show files debug-dump
- 4.5.6.1.13 show files stats
- 4.5.6.1.14 show files system
- 4.5.6.1.15 show files tcpdump

- 4.5.6.2 Configuration Files
  - 4.5.6.2.1 configuration audit
  - 4.5.6.2.2 configuration auto-upload
  - 4.5.6.2.3 configuration copy
  - 4.5.6.2.4 configuration delete
  - 4.5.6.2.5 configuration fetch
  - 4.5.6.2.6 configuration jump-start
  - 4.5.6.2.7 configuration merge
  - 4.5.6.2.8 configuration move
  - 4.5.6.2.9 configuration new
  - 4.5.6.2.10 configuration switch-to
  - 4.5.6.2.11 configuration text fetch
  - 4.5.6.2.12 configuration text file
  - 4.5.6.2.13 configuration text generate
  - 4.5.6.2.14 configuration upload
  - 4.5.6.2.15 configuration write
  - 4.5.6.2.16 write
  - 4.5.6.2.17 show configuration
  - 4.5.6.2.18 show configuration auto-upload
  - 4.5.6.2.19 show running-config
  - 4.5.6.2.20 show running-config interface

4.5.6.1 File System

4.5.6.1.1 debug generate dump

	debug generate dump Generates a debug dump.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Example	`switch (config) # debug generate dump` `Generated dump sysdump-switch-112104-201140526-091707.tgz`
Related Commands	file debug-dump
Notes	The dump can then be manipulated using the “file debug-dump...” commands.

4.5.6.1.2 file debug-dump

	file debug-dump {delete {<filename> | all | latest} | email {<filename> | latest} | upload {<filename> | latest | all [vrf <vrf-name>]} <URL>} Manipulates debug dump files.	
Syntax Description	delete	Deletes a debug dump file. • all—deletes all existing debug files from this machine • latest—deletes latest debug file from this machine
	email	Emails a debug dump file to pre-configured recipients for "informational events". • latest—emails the latest debug file to a pre-configured recipients
	upload	Uploads a debug dump file to a remote host. • latest—uploads the latest debug file to a remote host
	vrf-name—Describes VRF context that should be used for this transfer. If not specified, the "default" VRF is used.	
	URL	The URL to the remote host. Supported URL formats: HTTP, HTTPS, FTP, TFTP, SCP and SFTP. Example: scp://username[:password]@hostname/path/filename
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.3.4000 3.9.2000	Added "all" and "latest" options Added VRF option
Example	`switch (config) # file debug-dump email sysdump-switch-112104-20114052-091707.tgz`	
Related Commands	show files debug-dump	

4.5.6.1.3 file stats

	file stats {delete <filename> | move {<source filename> | <destination filename>} | upload <filename> [vrf <vrf-name>] <URL>} Manipulates statistics report files.	
Syntax Description	delete <filename>	Deletes a stats report file.
	move <source filename> <destination filename>	Renames a stats report file.

	upload <filename> <URL>	Uploads a stats report file. Supported URL formats: HTTP, HTTPS, FTP, TFTP, SCP and SFTP. Example: scp://username[:password]@hostname/path/filename
	vrf-name—Describes VRF context that should be used for this transfer. If not specified, the “default” VRF is used.	
Default	N/A	
Configuration Mode	config	
History	3.1.00003.9.2000—Added VRF option	
Example	`switch (config) # file stats move memory-1.csv memory-2.csv`	
Related Commands	show files stats show files stats <filename>	
Notes		

4.5.6.1.4 file tcpdump

	file tcpdump {delete <filename> | upload <filename> [vrf <vrf-name>] <URL>} Manipulates tcpdump output files.	
Syntax Description	delete <filename>	Deletes a stats report file.
	upload <filename> <URL>	Uploads the specified tcpdump output file to the specified URL. Supported URL formats: HTTP, HTTPS, FTP, TFTP, SCP and SFTP. Example: scp://username[:password]@hostname/path/filename.
	vrf-name—Describes VRF context that should be used for this transfer. If not specified, the “default” VRF is used.	
Default	N/A	
Configuration Mode	config	
History	3.1.0000 3.9.2000—Added VRF option	
Example	`switch (config) # file tcmpdump delete my-tcpdump-file.txt`	
Related Commands	show files stats tcpdump	
Notes		

4.5.6.1.5 file eula upload

	file eula upload <filename> <URL> Uploads the End User License Agreement to a specified remote location.	
Syntax Description	filename	The End User License Agreement
	URL	URL or scp://username[:password]@hostname/path/filename
Default	N/A	
Configuration Mode	config	

History	3.4.1100
Example	switch (config) # file eula upload MLNX-OS_EULA.pdf ? <URL or scp://username[:password]@hostname/path/filename>
Related Commands	license
Notes	N/A

4.5.6.1.6 file open-source-licenses upload

	file open-source-licenses upload <filename> <URL> Uploads the Open Source Licenses file.	
Syntax Description	filename	The Open Source Licenses file
	URL	URL or scp://username[:password]@hostname/path/filename
Default	N/A	
Configuration Mode	config	
History	3.9.3100	
Example	switch (config) # file open-source-licenses upload Open_Source_Licenses.txt scp://username[:password]@hostname/path/filename	
Related Commands	license	
Notes	N/A	

4.5.6.1.7 file help-docs upload

	file help-docs upload <filename> <URL or scp://username[:password]@hostname/path/filename> Uploads OS documentation to a specified remote location.	
Syntax Description	filename	The file to upload to a remote host.
	URL	URL or scp://username[:password]@hostname/path/filename.
Default	N/A	
Configuration Mode	config	
History	3.4.1100	
Example	switch (config) # file help-docs upload Onyx_ETH_User_Manual.pdf <scp://username[:password]@hostname/path/filename>	
Related Commands		
Notes		

4.5.6.1.8 reload

	reload [force immediate | halt [noconfirm] | noconfirm] Reboots or shuts down the system.

Syntax Description	force immediate	Forces an immediate reboot of the system even if the system is busy.
	halt	Shuts down the system.
	nonconfirm	Reboots the system without asking about unsaved changes.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	

Example	`switch (config) # reload`
	`Configuration has been modified; save first? [yes] yes`
	`Configuration changes saved.`
	`...`	

Related Commands	reset factory	
Notes		

4.5.6.1.9 reset factory

	reset factory [keep-all-config | keep-basic | keep-config-group | keep-virt-vols | keep-docker | keep-docker clear-label <label name>] | only-config] [halt] Clears the system and resets it entirely to its factory state.	
Syntax Description	keep-all-config	Preserves all configuration files including licenses. Removes the logs, stats, images, snapshots, history, and known hosts. The user is prompted for confirmation before honoring this command, unless confirmation is disabled with the command: “no cli default prompt confirm-reset”.
	keep-basic	Preserves licenses in the running configuration file.
	keep-config-group	Reset to the factory defaults of the current RoCE config group: no-roce, lossless, lossy or semi-lossless.
	keep-virt-vols	Preserves all virtual disk volumes.
	only-config	Removes configuration files only. Logs, stats, images, snapshots, history, and known hosts are preserved.
	halt	The system is halted after this process completes.
	keep-docker	Preserves all current docker configurations.
	keep-docker clear-label <label name>	Preserves all current docker configurations, but deletes the content of the given docker storage label. (Note that only the content of the label folder will be deleted. The label itself will remain intact.)
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Added notes and “keep-virt-vols” parameter
	3.6.2002	Updated example and notes
	3.8.1300	Added "keep-docker" and "keep-docker clear-label" option

Example	switch (config) # reset factory Warning - confirming will cause system reboot. Type 'YES' to confirm reset: YES Resetting and rebooting the system -- please wait... ...
Related Commands	reload
Notes	• Effects of parameter “keep-all-config”: Licenses—not deleted; profile—no change; configuration—unchanged; management IP—unchanged • Effects of parameter “keep-basic”: Licenses—not deleted; profile—reset; configuration—reset; management IP—reset • Effects of parameter “keep-virt-vols”: Licenses—deleted; profile—reset; configuration—reset; management IP—deleted • Confirming the command causes system reboot

4.5.6.1.10 configuration new factory

	configuration new <filename> factory Creates new file with only factory defaults.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.7.1102
Example	switch (config) # no configuration new my_file factory
Related Commands	configuration new factory configuration new factory keep-basic configuration new factory keep-connect
Notes	

4.5.6.1.11 configuration new factory keep-docker

	configuration new <filename> factory keep-docker Creates new file with only factory defaults except docker current configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.7.1102
Example	switch (config) # no configuration new my_file factory keep-docker
Related Commands	configuration new factory configuration new factory keep-basic configuration new factory keep-connect
Notes	

4.5.6.1.12 show files debug-dump

	show files debug-dump [<filename>] Displays a list of debug dump files.	
Syntax Description	filename	Displays a summary of the contents of a particular debug dump file.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example	switch (config) # show files debug-dump sysdump-switch-20170731-161038.tgz == System information: Hostname: switch Version: X86_64 3.6.4006 2017-07-03 16:17:39 x86_64 Current time: 2017-07-31 16:10:38 System uptime: 19d 18h 20m 12s == == Output of 'uname -a': Linux switch 3.10.0-327.36.3.el7smp-x86_64 X86_64 jenkins #1 2017-06-27 12:34:55 SMP x86_64 x86_64 x86_64 GNU/Linux ==	
Related Commands	file debug-dump	
Notes		

4.5.6.1.13 show files stats

	show files stats <filename> Displays a list of statistics report files.	
Syntax Description	filename	Display the contents of a particular statistics report file.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example	switch (config) # show files stats memory-201140524-111745.csv	
Related Commands	file stats	
Notes		

4.5.6.1.14 show files system

	show files system [detail] Displays usage information of the file systems on the system.

Syntax Description	detail	Displays more detailed information on file-system.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example	switch (config) # show files stats memory-201140524-111745.csv	
Related Commands		
Notes		

4.5.6.1.15 show files tcpdump

	show files tcpdump Displays a list of statistics report files.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config) # show files stats test dump3
Related Commands	
Notes	

4.5.6.2 Configuration Files

4.5.6.2.1 configuration audit

	configuration audit max-changes <number> Chooses settings related to configuration change auditing.	
Syntax Description	max-changes	Set maximum number of audit messages to log per change.
Default	1000	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # configuration audit max-changes 100	
Related Commands	show configuration	
Notes		

4.5.6.2.2 configuration auto-upload

	configuration auto-upload remote-url no configuration auto-upload remote-url Sets the remote URL to upload for automated backup. The no form resets the remote URL.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.9.0500
Example	`switch (config) # configuration auto-upload remote-url "scp://root:password@192.168.10.125/tmp/conf1"`
Related Commands	show configuration auto-upload
Notes	If this feature is set, after every configuration write it will upload the active configuration file to the configured remote URL.

4.5.6.2.3 configuration copy

	configuration copy <source-name> <dest-name> Copies a configuration file.	
Syntax Description	source-name	Name of source file.
	dest-name	Name of destination file. If the file of specified filename does not exist a new file will be created with said filename.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # configuration copy initial.bak example`	
Related Commands		
Notes	• This command does not affect the current running configuration • The active configuration file may not be the target of a copy. However, it may be the source of a copy in which case the original remains active.	

4.5.6.2.4 configuration delete

	configuration delete <filename> Deletes a configuration file.	
Syntax Description	filename	Name of file to delete
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # configuration delete example`	

Related Commands	show configuration files
Notes	• This command does not affect the current running configuration • The active configuration file may not be deleted

4.5.6.2.5 configuration fetch

	configuration fetch <URL> [<name>] Downloads a configuration file from a remote host.	
Syntax Description	URL	Supported formats: HTTP, HTTPS, FTP, TFTP, SCP and SFTP. Example: scp://username[:password]@hostname/path/filename
	name	The name of the configuration file.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # configuration fetch scp://root:password@192.168.10.125/tmp/conf1`	
Related Commands	configuration switch-to	
Notes	• The downloaded file should not override the active configuration file, using the <name> parameter • If no name is specified for a configuration fetch, it is given the same name as it had on the server • No configuration file may have the name “active”	

4.5.6.2.6 configuration jump-start

	configuration jump-start Runs the initial-configuration wizard.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000

Example	`switch (config) # configuration jump-start`
	`Configuration wizard`
	`Step 1: Hostname? [switch-3cc29c]`
	`Step 2: Use DHCP on mgmt0 interface? y`
	`Step 3: Admin password (Enter to leave unchanged)?`
	`You have entered the following information:`
	`1. Hostname: switch-3cc29c`
	`2. Use DHCP on mgmt0 interface: yes`
	`3. Enable IPv6: yes`
	`4. Enable IPv6 autoconfig (SLAAC) on mgmt0 interface: yes`
	`53. Admin password (Enter to leave unchanged): (unchanged)`
	`To change an answer, enter the step number to return to.`
	`Otherwise hit <enter> to save changes and exit.`
	`Choice:`
	`Configuration changes saved.`

Related Commands	configuration switch-to

Notes	• The wizard is automatically invoked whenever the CLI is launched when the active configuration file is fresh (i.e. not modified from its initial contents) • This command invokes the wizard on demand (see "Configuring the Switch for the First Time")

4.5.6.2.7 configuration merge

	configuration merge <filename> Merges the "shared configuration" from one configuration file into the running configuration.	
Syntax Description	filename	Name of file from which to merge settings.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # configuration merge new-config-file`	
Related Commands		
Notes	• No configuration files are modified during this process • The configuration filename must be a non-active configuration file	

4.5.6.2.8 configuration move

	configuration move <source-name> <dest-name> Renames a configuration file.	
Syntax Description	source-name	Name of file to rename.
	dest-name	New name of renamed file.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	

Example	`switch (config) # show configuration files`
	`example1 initial initial.bak initial.prev`
	`switch (config) # configuration move example1 example2`
	`switch (config) # show configuration files`
	`example2 initial initial.bak initial.prev`	

Related Commands	show configuration	
Notes	• This command does not affect the current running configuration • The active configuration file may not be the target of a move	

4.5.6.2.9 configuration new

	configuration new <filename> [factory [keep-basic] [keep-connect]] Creates a new configuration file under the specified name. The parameters specify what configuration, if any, to carry forward from the current running configuration.	
Syntax Description	filename	Names for new configuration file.
	factory	Creates new file with only factory defaults.

	keep-basic	Keeps licenses and host keys.
	keep-connect	Keeps configuration necessary for connectivity (interfaces, routes, and ARP).
Default	Keeps licenses and host keys	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # show configuration files initial initial.bak initial.prev switch (config) # configuration new example2 switch (config) # show configuration files example2 initial initial.bak initial.prev	
Related Commands	show configuration	
Notes	• This command does not affect the current running configuration • The active configuration file may not be the target of a move	

4.5.6.2.10 configuration switch-to

	configuration switch-to <filename>[no-reboot] Loads the configuration from the specified file and makes it the active configuration file.
Syntax Description	no-reboot | Forces configuration change without rebooting.
Default	N/A
Configuration Mode	config
History	3.1.0000 3.6.1002 | Added “no-reboot” option
Example	switch (config) # show configuration files initial (active) newcon initial.prev initial.bak switch (config) # configuration switch-to newcon no-reboot switch (config) # show configuration files initial newcon (active) initial.prev initial.bak
Related Commands	show configuration files
Notes	• The current running configuration is lost and not automatically saved to the previous active configuration file • When running the command without the “no-reboot” parameter, the user is prompted to OK a reboot. If the answer is “yes”, the configuration is replaced and the system is rebooted immediately

4.5.6.2.11 configuration text fetch

	configuration text fetch <URL> [apply [discard | fail-continue | filename | overwrite | verbose] | filename <filename> | overwrite [apply | filename <filename>]] Fetches a text configuration file (list of CLI commands) from a specified URL.

Syntax Description	apply	Applies the file to the running configuration (i.e. executes the commands in it). This option has the following parameters: • discard—does not keep downloaded configuration text file after applying it to the system • fail-continue—if applying commands, continues execution even if one of them fails • overwrite—if saving the file and the filename already exists, replaces the old file • verbose—displays all commands being executed and their output instead of just those that get errors
	filename	Specifies filename for saving downloaded text file.
	overwrite	Downloads the file and saves it using the same name it had on the server. This option has the following parameters: • apply—applies the downloaded configuration to the running system • filename—specifies filename for saving downloaded text file
Default	N/A	
Configuration Mode	config	
History	3.2.1000	
Example	`switch (config) # configuration text fetch scp://username[:password]@hostname/path/filename`	
Related Commands		
Notes		

4.5.6.2.12 configuration text file

	configuration text file <filename> {apply [fail-continue] [verbose] [reboot] | delete | rename <filename> | upload < URL>} Performs operations on text-based configuration files.	
Syntax Description	filename <file>	Specifies the filename.
	apply	Applies the configuration on the system.
	fail-continue	Continues execution of the commands even if some commands fail.
	verbose	Displays all commands being executed and their output, instead of just those that get errors.
	delete	Deletes the file.
	rename <filename>	Renames the file.
	upload <URL>	Supported types are HTTP, HTTPS, FTP, TFTP, SCP and SFTP. For example: scp://username[:password]@hostname/path/filename
	reboot	Write the configuration and reboot after successful execution.
Default	N/A	
Configuration Mode	config	
History	3.1.0000 3.9.0300 | Added ability to apply reboot	

Example	`switch (config) # configuration text file my-config-file delete`
Related Commands	show configuration files
Notes	

4.5.6.2.13 configuration text generate

	configuration text generate {active {running | saved} | file <filename> } {save <filename> | upload <URL>} Generates a new text-based configuration file from this system's configuration.	
Syntax Description	active	Generates from currently active configuration.
	running	Uses running configuration.
	saved	Uses saved configuration.
	file <filename>	Generates from inactive saved configuration.
	save	Saves new file to local persistent storage.
	upload <URL>	Supported types are HTTP, HTTPS, FTP, TFTP, SCP and SFTP. For example: scp://username[:password]@hostname/path/filename.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # configuration text generate file initial.prev save example`	
Related Commands	show configuration files	
Notes		

4.5.6.2.14 configuration upload

	configuration upload {timestamp} {active | <name>} <URL or scp or sftp://username:password@hostname[:port]/path/filename> Uploads a configuration file to a remote host.	
Syntax Description	active	Upload the active configuration file.
	timestamp	Will append the timestamp to the filename uploaded to remote.
Default	N/A	
Configuration Mode	config	
History	3.1.0000 3.9.0500 | Added timestamp option	
Example	`switch (config) # configuration upload active scp://root:password@192.168.10.125/tmp/conf1`	
Related Commands	show configuration files	
Notes	No configuration file may have the name “active” or “timestamp”.	

4.5.6.2.15 configuration write

	configuration write [local | to <filename> [no-switch]] Saves the running configuration to the active configuration file.	
Syntax Description	local	Saves the running configuration locally (same as “write memory local”).
	to <filename>	Saves the running configuration to a new file under a different name and makes it the active file.
	no-switch	Saves the running configuration to this file but keep the current one active.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # configuration write`	
Related Commands	write	
Notes		

4.5.6.2.16 write

	write {memory [local] | terminal} Saves or displays the running configuration.	
Syntax Description	memory	Saves running configuration to the active configuration file. It is the same as “configuration write”.
	local	Saves the running configuration only on the local node. It is the same as “configuration write local”.
	terminal	Displays commands to recreate current running configuration. It is the same as “show running-config”.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	

Example	switch (config) # write terminal ## ## Running database "initial" ## Generated at 2014/05/27 10:05:16 +0000 ## Hostname: switch ## ## ## Network interface configuration ## interface mgmt0 comment "" interface mgmt0 create interface mgmt0 dhcp interface mgmt0 display interface mgmt0 duplex auto interface mgmt0 mtu 1500 no interface mgmt0 shutdown interface mgmt0 speed auto no interface mgmt0 zeroconf ## ## Local user account configuration ## username a** capability admin no username a** disable username a** disable password .....
Related Commands	show running-config configuration write
Notes	

4.5.6.2.17 show configuration

	show configuration [audit | files [<filename>] | running | text files] Displays a list of CLI commands that will bring the state of a fresh system up to match the current persistent state of this system.	
Syntax Description	audit	Displays settings for configuration change auditing.
	files [<filename>]	Displays a list of configuration files in persistent storage if no filename is specified. If a filename is specified, it displays the commands to recreate the configuration in that file. In the latter case, only non-default commands are shown, as for the normal “show configuration” command.
	running	Displays commands to recreate current running configuration. Same as the command “show configuration” except that it applies to the currently running configuration, rather than the current persisted configuration.
	text files	Displays names of available text-based configuration files.
Default	N/A	
Configuration Mode	config	
History	3.1.0000 3.3.5006 | Removed “running full” and “full” parameters	

Example	switch (config) # show configuration ## ## Active saved database "newcon" ## Generated at 20114/05/25 10:18:52 +0000 ## Hostname: switch-3cc29c ## ## ## Network interface configuration ## interface mgmt0 comment "" interface mgmt0 create interface mgmt0 dhcp interface mgmt0 display interface mgmt0 duplex auto interface mgmt0 mtu 1500 no interface mgmt0 shutdown interface mgmt0 speed auto no interface mgmt0 zeroconf
Related Commands	
Notes	

4.5.6.2.18 show configuration auto-upload

	show configuration auto-upload Shows the automated backup settings.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.9.0500
Example	switch (config) # show configuration auto-upload Auto-upload settings: Enabled: yes Remote url: scp://root@192.168.10.125/tmp/conf1 Password : ******
Related Commands	configuration auto-upload remote-url
Notes	If this feature is set. After every configuration write, it will upload the active configuration file to the configured remote URL.

4.5.6.2.19 show running-config

	show running-config [expanded | protocol <protocol>| diff | diff <config_file_name>] Displays commands to recreate current running configuration.	
Syntax Description	expanded	Displays commands in expanded format without compressing ranges.
	protocol	Only displays commands relating to the specified protocol.
	diff	Displays delta between saved config file (active by default) and running-config.
	config_file_name	Displays delta between the specified saved config file and running-config.
Default	N/A	

Configuration Mode	config	
History	3.1.0000	
	3.3.4402	Removed “full” parameter
	3.6.2002	Updated example and added parameters
	3.6.3640	Added support for forwarding mode configuration
	3.8.1000	Added support to show diff between running-config and saved config files (active file saved by default)
Example		

```
switch (config) # show running-config diff

Only in running-config:
+ interface port-channel 1
+ interface ethernet 1/31-1/33 speed 10G force
+ interface port-channel 1 description lag
Only in saved configuration file:
- ip route vrf default 169.254.22.0/24 169.254.2.100
Common configuration but in different order in saved configuration file and running-config:
<<None>>
```

Related Commands	
Notes	• + <string> : <string> exists only in running-config, but not in the saved filename (or active config file if no <filename> is specified) • - <string> : <string> does not exist in running-config, but exists in the saved filename (or active config file if no <filename> is specified) • ! <string> : <string> exists in both running-config and the saved filename, but it is out of order. This should not impact the user, but may impact scripts or applications that are parsing the output of the command

4.5.6.2.20 show running-config interface

	show running-config interface [mgmt0 | mgmt1 | lo <loopback_id> | ethernet <slot>/<port>[/<subport>] | port-channel <lag-id> | mlag-port-channel <mlag-id> | nve <nve-id> | vlan <vlan-id>] Displays running-config filtered with the specific interfaces.	
Syntax Description	loopback_id	Loopback interface ID. Range: 0-31
	<slot>/<port>	Ethernet port number.
	subport	Ethernet subport number
	lag-id	LAG ID number. Range: 1-4096
	mlag-id	MLAG ID number. Range: 1-1000
	nve-id	NVE ID number. Range: 1-64

	vlan-id	VLAN ID number. Range: 1-4094
Default	N/A	
Configuration Mode	config	
History	3.8.1000	
Example	switch (config) # show running-config interface mgmt0 interface mgmt0 comment mgmt if switch (config) # show running-config interface mgmt1 interface mgmt1 comment mgmt if switch (config) # show running-config interface lo 1 interface loopback 1 interface loopback 1 ip address 1.1.10.10/32 primary switch (config) # show running-config interface ethernet 1/32 interface ethernet 1/32 speed 10G force switch (config) # show running-config interface port-channel 1 interface port-channel 1 interface port-channel 1 description lag switch (config) # show running-config interface mlag-port-channel 1 interface mlag-port-channel 1 interface mlag-port-channel 1 description mlag switch (config) # show running-config interface nve 1 interface nve 1 interface nve 1 nve fdb learning remote interface nve 1 nve fdb flood load-balance switch (config) # show running-config interface vlan 100 interface vlan 100 interface vlan 100 ip address 169.254.1.101/24 primary interface vlan 100 ip address 169.254.11.101/24	
Related Commands		
Notes		

4.6 Resource Scale

NVIDIA Onyx allows dynamic allocation of internal resources so that different internal subsystems could use as much resources as are available until resource exhaustion is reached.

Internal subsystems (e.g., ACL, OF, IP router) may use internal resources according to configured allocation policy mode which, in the case of Spectrum-based switch systems is loose. Loose mode is a configuration that supports flexible user experience while providing protection to assure some protection against flooding of ARP.

> Transition between modes saves configuration and reloads the system.

The following table presents the number of resources available for a NVIDIA Spectrum™Spectrum™-based node in loose mode.

4.6.1 Resource Scale Commands

4.6.1.1 show system resource table

	show system resource table [<table-id>] Displays all system resource in-use value.	
Syntax Description	table-id	Displays information for a specific in-use resource table

Default	N/A
Configuration Mode	Any command mode
History	3.5.1000

Example	`switch (config) # show system resource table`
	`---------------------------------`
	`Table-Id In-Use`
	`---------------------------------`
	`acl 0`
	`ipv4-uc 1`
	`ipv4-mc 0`
	`ipv4-neigh 0`
	`ipv6-uc 0`
	`ipv6-mc 0`
	`ipv6-neigh 0`
	
	`System mode: loose`
	`Total configured entries: 1`

Related Commands	
Notes	

5 System Synchronization

The following pages provide information on NTP and PTP functionalities.

- NTP and Clock
- Precision Time Protocol (PTP)
- Replace CRC with Timestamp

5.1 NTP and Clock

Network Time Protocol (NTP) is a networking protocol for clock synchronization between computer systems over variable-latency data networks. NTP is intended to synchronize all participating computers to within a few milliseconds of Coordinated Universal Time (UTC) and is designed to mitigate the effects of variable network latency. NTP can usually maintain time to within tens of milliseconds over the public Internet, and can achieve better than one millisecond accuracy in local area networks under ideal conditions.

5.1.1 NTP Authenticate

When authentication of incoming NTP packets is enabled, the switch ensures that they come from an authenticated time source before using them for time synchronization on the switch. Authentication keys are created and added to the trusted list.

To add a key to be used for authentication, take the following steps.

1. Create the key.

```
switch (config)# ntp authentication-key 1 md5 password
```

2. Add the key to the trusted list.

```
switch (config)# ntp trusted-key 1
```

3. Assign the key to the server/peer.

```
switch (config)# ntp server 10.34.1.1 keyID 1
```

5.1.2 NTP Authentication Key

An authentication key may be created and used to authenticate incoming NTP packets. For the key to be used, make sure the following is in place.

1. It should be shared with the NTP server/peer sending the NTP packet.
2. It should be added to the trusted list.
3. NTP authenticate should be enabled on the system

5.1.3 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community post:

- HowTo Enable NTP

5.1.4 NTP Commands

- 5.1.1 NTP Authenticate
- 5.1.2 NTP Authentication Key
- 5.1.3 Additional Reading and Use Cases
- 5.1.4 NTP Commands
  - 5.1.4.1 clock set
  - 5.1.4.2 clock timezone
  - 5.1.4.3 ntp
  - 5.1.4.4 ntpdate
  - 5.1.4.5 ntp authenticate
  - 5.1.4.6 ntp authentication-key
  - 5.1.4.7 ntp peer disable
  - 5.1.4.8 ntp peer keyID
  - 5.1.4.9 ntp peer version
  - 5.1.4.10 ntp server disable
  - 5.1.4.11 ntp server keyID
  - 5.1.4.12 ntp server-role disable
  - 5.1.4.13 ntp server trusted-enable
  - 5.1.4.14 ntp server version
  - 5.1.4.15 ntp trusted-key
  - 5.1.4.16 show clock
  - 5.1.4.17 show ntp
  - 5.1.4.18 show ntp configured
  - 5.1.4.19 show ntp keys

5.1.4.1 clock set

1		clock set <hh:mm:ss> [<yyyy/mm/dd>] Sets the time and date.	
2	Syntax Description	hh:mm:ss	Time
		yyyy/mm/dd	Date
3	Default	N/A	
4	Configuration Mode	config	
5	History	3.1.0000	

6	Example	`switch (config) # clock set 23:23:23 2010/08/19`
7	Related Commands	show clock
8	Notes	If not specified, the date will be left the same.

5.1.4.2 clock timezone

	clock timezone [<zone-word> [<zone-word> [<zone-word>] [<zone-word>]]] no clock timezone Sets the system time zone. The time zone may be specified in one of three ways: • A nearby city whose time zone rules to follow. The system has a large list of cities which can be displayed by the help and completion system. They are organized hierarchically because there are too many of them to display in a flat list. A given city may be required to be specified in two, three, or four words, depending on the city • An offset from UTC. This will be in the form UTC-offset UTC, UTC-offset UTC+<0-14>, UTC-offset UTC-<1-12> • UTC (Universal Time, which is almost identical to GMT), and this is the default time zone The no form of the command resets time zone to its default (GMT).	
Syntax Description	zone-word	Possible forms this could take include: continent, city, continent, country, city, continent, region, country, city, ocean, and/or island.
Default	GMT	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # clock timezone America North United_States Other New_York`	
Related Commands	show clock	
Notes		

5.1.4.3 ntp

	ntp { {[vrf <vrf-name>] { disable | enable [force]}} | {peer | server} <IP address> [version <number> | disable]} no ntp { {[vrf <vrf-name>] {disable | enable}} | {peer | server} <IP address> [version <number> | disable]} Configures NTP. The no form of the command negates NTP options.	
Syntax Description	disable	Disables NTP.
	enable	Enables NTP.
	peer | server	Configures an NTP peer or server node.
	IP address	IPv4 or IPv6 address.
	version <number>	Specifies the NTP version number of this peer. Possible values: 3 or 4

	vrf name—Describes the VRF name for NTP daemon. If the VRF parameter is not specified, the "default" VRF will be used implicitly force—This option will restart ntp with passed VRF context even if it was already enabled using other VRF.
Default	NTP is enabled NTP version number is 4
Configuration Mode	config
History	3.1.0000 3.9.2000—Added VRF option
Example	`switch (config) # no ntp peer 192.168.10.24 disable`
Related Commands	
Notes	NTP can be enabled only in one VRF at a time.

5.1.4.4 ntpdate

	ntpdate <ip-address> Configures the system clock using the specified SNTP server.	
Syntax Description	ip-address	IP address of SNTP server.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # ntpdate 192.168.10.10` `26 Feb 17:25:40 ntpdate[15206]: adjust time server 192.168.10.10 offset -0.000092 sec`	
Related Commands		
Notes	This is a one-time operation and does not cause the clock to be kept in sync on an ongoing basis. It will generate an error if SNTP is enabled since the socket it requires will already be in use.	

5.1.4.5 ntp authenticate

	ntp authenticate no ntp authenticate Enables NTP authentication. The no form of the command disables NTP authentication.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.5.0200
Example	`switch (config) # ntp authenticate`
Related Commands	
Notes	

5.1.4.6 ntp authentication-key

	ntp authentication-key <key-id> <encrypt-type> [<password>] no ntp authentication-key <key-id> Enables NTP authentication. The no form of the command disables NTP authentication.	
Syntax Description	key-id	Specifies a key ID, whether existing or a new one to be added. Range: 1-65534
	encrypt-type	Specifies encryption type to use (md5, or sha1)
	password	Password string
Default	Disabled	
Configuration Mode	config	
History	3.5.0200	

Example	`switch (config) # ntp authentication-key 123 md5 examplepass`
	`switch (config) # ntp authentication-key 1234 sha1`
	`Password: **`
	`Confirm: **`	

Related Commands		
Notes	If a password is not entered, a prompt appears requiring that a password is introduced.	

5.1.4.7 ntp peer disable

	ntp peer <ip-address> disable no ntp peer <ip-address> disable Temporarily disables this NTP peer. The no form of the command enables this NTP peer.	
Syntax Description	ip-address	IP address of the peer.IPv4, IPv6 and hostname (FQDN) are acceptable.
Default	Disabled	
Configuration Mode	config	
History	3.5.0200 3.6.4000—Added hostname as option for ip-address, and added note	
Example	`switch (config) # ntp peer 10.10.10.10 disable`	
Related Commands		
Notes	• IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone ID for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7'.) • The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen.	

5.1.4.8 ntp peer keyID

	ntp peer <ip-address> keyID <key-id> no ntp peer <ip-address> keyID <key-id> Specifies the KeyID of the NTP peer. The no form of the command removes key ID configuration from the NTP peer.	
Syntax Description	ip-address	IP address of the peer.IPv4, IPv6 and hostname (FQDN) are acceptable.
	key-id	Range: 1-65534
Default	Disabled	
Configuration Mode	config	
History	3.5.0200 3.6.4000—Added hostname as option for ip-address, and added note	
Example	`switch (config) # ntp peer 10.10.10.10 keyID 120`	
Related Commands		
Notes	• IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone ID for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7'.) • The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen.	

5.1.4.9 ntp peer version

	ntp peer <ip-address> version <ver-num> no ntp peer <ip-address> version <ver-num> Specifies the NTP version number of this peer. The no form of the command defaults NTP to version 4.	
Syntax Description	ip-address	IP address of the peer.IPv4, IPv6 and hostname (FQDN) are acceptable.
	ver-num	NTP version. Possible values: 3 or 4
Default	4	
Configuration Mode	config	
History	3.5.0200 3.6.4000—Added hostname as option for ip-address, and added note	
Example	`switch (config) # ntp peer 10.10.10.10 version 4`	
Related Commands		
Notes	• IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone ID for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7') • The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen.	

5.1.4.10 ntp server disable

	ntp server <ip-address> disable no ntp server <ip-address> disable Temporarily disables this NTP server. The no form of the command enables this NTP server.	
Syntax Description	ip-address	IP address of the peer.IPv4, IPv6 and hostname (FQDN) are acceptable.
Default	Disabled	
Configuration Mode	config	
History	3.5.5000 3.6.4000—Added hostname as option for ip-address, and added note	
Example	`switch (config) # ntp server 10.10.10.10 disable`	
Related Commands		
Notes	• IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone ID for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7'.) • The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen.	

5.1.4.11 ntp server keyID

	ntp server <ip-address> keyID <key-id> no ntp server <ip-address> keyID <key-id> Specifies the KeyID of the NTP server. The no form of the command removes key ID configuration from the NTP server.	
Syntax Description	ip-address	IP address of the peer.IPv4, IPv6 and hostname (FQDN) are acceptable.
	key-id	Range: 1-65534
Default	Disabled	
Configuration Mode	config	
History	3.5.0200 3.6.4000—Added hostname as option for ip-address, and added note	
Example	`switch (config) # ntp server 10.10.10.10 keyID 120`	
Related Commands		
Notes	• IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone ID for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7'.) • The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen.	

5.1.4.12 ntp server-role disable

	ntp server-role disable no ntp server-role disable Disables the switch's default ability to function as an NTP server. The no form of the command restores the switch's ability to function as an NTP server.
Syntax Description	N/A
Default	N/A
Configuration Mode	Configure terminal
History	3.8.2100
Role	Admin
Example	`switch (config) # ntp server-role disable`
Related Commands	show ntp
Notes	This command is configurable.

5.1.4.13 ntp server trusted-enable

	ntp server <ip-address> trusted-enable no ntp server <ip-address> trusted-enable Trusts this NTP server; if authentication is configured this will additionally force all time updates to only use trusted servers. The no form of the command removes trust from this NTP server.	
Syntax Description	ip-address	IP address of the peer.IPv4, IPv6 and hostname (FQDN) are acceptable.
Default	N/A	
Configuration Mode	config	
History	3.6.2002 3.6.4000—Added hostname as option for ip-address, and added note	
Example	`switch (config) # ntp server 10.10.10.10 trusted-enable`	
Related Commands		
Notes	• IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone ID for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7'.) • The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen. • NTP trusted servers can be used as a mitigation for Sybil attacks which is a vulnerability caused by NTP peers sharing the same NTP key base. This mitigation adds the concept of trusted servers which if enabled in conjunction with NTP authentication ensures that time information will only be obtained from trusted servers.	

5.1.4.14 ntp server version

	ntp server <ip-address> version <ver-num> no ntp server <ip-address> version <ver-num> Specifies the NTP version number of this server. The no form of the command defaults NTP to version 4.	
Syntax Description	ip-address	IP address of the peer.IPv4, IPv6 and hostname (FQDN) are acceptable.
	ver-num	NTP version. Possible values: 3 or 4
Default	4	
Configuration Mode	config	
History	3.5.0200 3.6.4000—Added hostname as option for ip-address, and added note	
Example	`switch (config) # ntp server 10.10.10.10 version 4`	
Related Commands		
Notes	• IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone ID for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7') • The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen	

5.1.4.15 ntp trusted-key

	ntp trusted-key <key(s)> no ntp trusted-key <key(s)> Adds one or more keys to the trusted key list. The no form of the command removes keys from the trusted key list.	
Syntax Description	key(s)	Range: 1-65534
Default	Disabled	
Configuration Mode	config	
History	3.5.0200	
Example	`switch (config) # ntp trusted-key 1,3,5` `switch (config) # ntp trusted-key 1-5`	
Related Commands		
Notes	Keys may be separated with commas without any space, or they may be set as a range using a hyphen.	

5.1.4.16 show clock

	show clock Displays the current system time, date and time zone.
Syntax Description	N/A
Default	N/A

Configuration Mode	Any command mode
History	3.1.0000 3.6.6000—Updated example

Example	`switch (config) # show clock`
	
	`Time: 02:48:41`
	`Date: 2018/1/1`
	`Time zone: UTC (Etc/UTC)`
	`UTC offset: same as UTC`

Related Commands	
Notes	

5.1.4.17 show ntp

	show ntp Displays the current NTP settings.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.5.0200—Updated example 3.6.6000—Updated example 3.9.2000—Updated example

Example	`switch (config)# show ntp`
	
	`NTP is administratively : enabled`
	`VRF name : mgmt`
	`NTP Authentication administratively: disabled`
	`NTP server role : enabled`
	
	`Clock is synchronized:`
	`Reference: 10.7.7.134`
	`Offset : -0.038 ms`
	
	`Active servers and peers:`
	`10.7.7.134:`
	`Conf Type : serv`
	`Status : sys.peer(*)`
	`Stratum : 3`
	`Offset(msec) : -0.038`
	`Ref clock : 192.14.55.225`
	`Poll Interval (sec): 128`
	`Last Response (sec): 101`
	`Auth state : none`

Related Commands	
Notes	

5.1.4.18 show ntp configured

	show ntp configured Displays NTP configuration.
Syntax Description	N/A

Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.6.6102—Updated example
Example	

```
switch (config) # show ntp configured


NTP enabled: yes
NTP Authentication enabled: no
NTP peer 0.us.pool.ntp.org # Hostname peer configuration
 Resolved as: 45.79.111.114
 Enabled: yes
 NTP version: 4
 Key ID: none
NTP peer 2.3.1.3 # IP peer configuration
 Enabled: yes
 NTP version: 4
 Key ID: none
NTP server vnc23 # Hostname server configuration
 Resolved as: 10.7.2.23
 Enabled: yes
 NTP version: 4
 Key ID: none
 Trusted: no
NTP server 1.2.3.4 # IP server configuration
 Enabled: yes
 NTP version: 4
 Key ID: none
 Trusted: no
NTP server idontexist (DNS resolution failed. Reset or reconfigure NTP to try again)
 Enabled: yes
 NTP version: 4
 Key ID: none
 Trusted: no
```

Related Commands	
Notes	

5.1.4.19 show ntp keys

	show ntp configured Displays NTP keys.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.5.0200

Example	`switch (config) # show ntp keys`
	`NTP Key 1`
	`  Trusted: yes`
	`  Encryption Type: MD5`
	`NTP Key 2`
	`  Trusted: yes`
	`  Encryption Type: MD5`
	`NTP Key 3`
	`  Trusted: yes`
	`  Encryption Type: MD5`
	`NTP Key 4`
	`  Trusted: yes`
	`  Encryption Type: md5`

Related Commands	

Notes	

5.2 Precision Time Protocol (PTP)

Synchronizing network applications require their wall clock time to be aligned precisely with a reference time source (to the order of micro seconds or less). To achieve such accuracy, the application needs the support of networking HW (switch and adapter card), to provide the means to stamp time-sensitive packets. It also requires a time synchronization protocol which would make use of the HW time stamping to adjust its wall clock time to an accurate clock in the network.

5.2.1 PTP Principles

The basic principle of PTP is as follows: Slave time = master time + propagation delay + offset.

The purpose of the protocol is to align the slave and the master time so that the gap between them is the propagation delay of the packet. Or in other words, the purpose of the protocol is to use the offset to correct the slave time so the offset between the master sending the packet and the slave receiving the packet is the propagation delay.

Master time is sent periodically by a reliable clock source named Master Clock (MC). In a PTP network, one single reference source is elected called Grand Master Clock (GMC). Propagation delay is calculated between each node and the MC by one of the two methods provided by the standard and further explained below.

To reach sub-microsecond resolutions, all the time stamps which record when a packet is sent and received should be done in the HW. This may impose interaction between SW and HW to query the HW time and send follow-up messages. This issue is further explained below in 2 step section.

Assuming that the propagation delay in the network is symmetric, the propagation time is the average time that took the sync and delay req messages to be switched.

Propagation delay = (T4-T1-(T3-T2))/2=(T4-T1+T2-T3)/2

T1 represents the time that the packet left the master which is actually the master time.

The following figure provides an example of the stages required by a slave clock to align its time to the master clock:

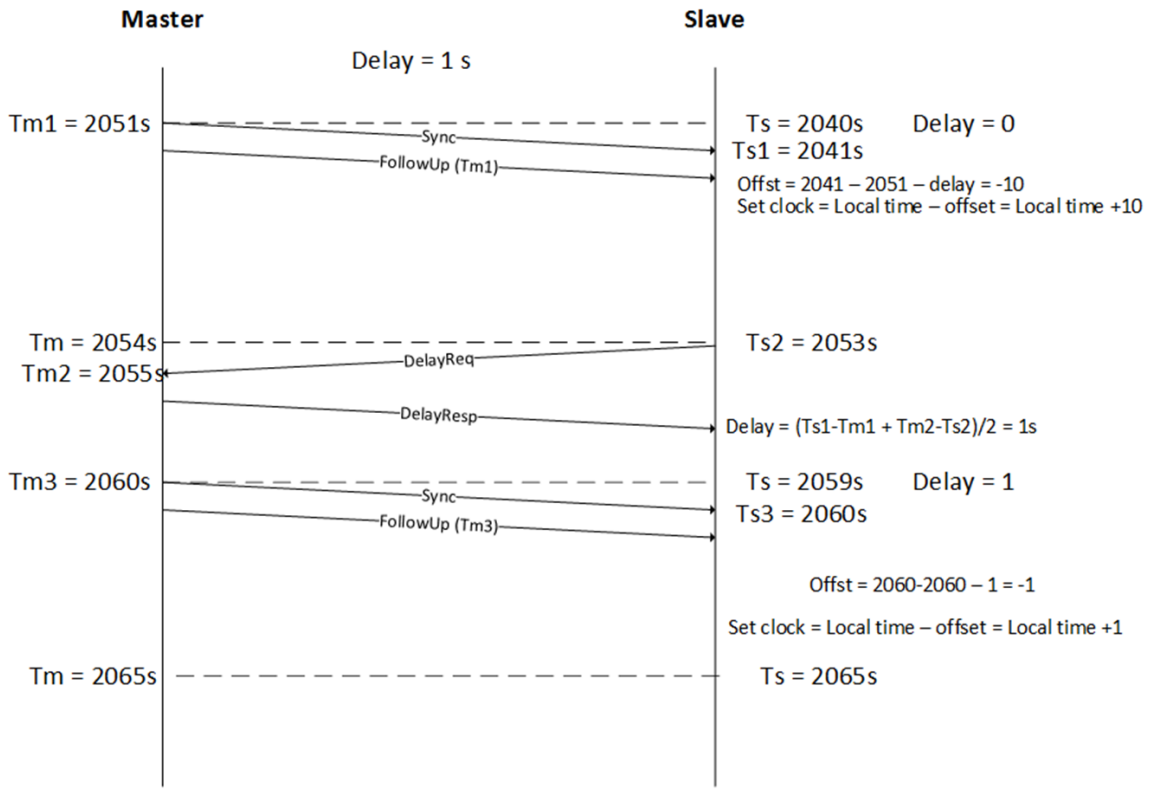

The following table presents the PTP message formats:

Message Type	Hex Value	Class
Sync	0	Event
Follow-up	8	General
Delay_Req	1	Event
Delay_Resp	9	General
Pdelay_Req	2	Event
Pdelay_Resp	3	Event
Pdelay_Resp follow-up	A	General
Announce	B	General
Signaling	C	General
Management	D	General

5.2.2 Clock Types and Operation Modes

The types of clocks available are as follows:

- Grand Master Clock (GMC)—the reference time source derived from an accurate clock such as a GNSS driven clock (i.e. GPS, GLONASS, GALILEO)
- Boundary Clock (BC)—a network device that acts as slave to its master and as master to its slaves. (NVIDIA Onyx implements only this)

- Ordinary Clock (OC)—a clock that operates either as a Master or a Slave. In the case of a slave, the end point whose clock is been synced (normally a host/server).
- Master Clock (MC)—a clock which operates as a Master and derives its timing capabilities from the clock chain up to the GMC. It typically serves as a port on a BC connected to a host running as a slave.
- Transparent Clock (TC)—a PTP aware switch capable of measuring the PTP packet switching delay (transient time) and updating the data in the packet. In peer-to-peer (P2P) delay calculation mechanism, a TC device is also required to calculate its delay from the next hop toward the MC and add the value to the switching delay.

Two modes of delay calculations are defined:

- End-to-End (E2E)—each slave calculates its delay from the MC by running Delay request/ delay response sequence (NVIDIA Onyx implements only this)
- Peer-to-Peer—propagation delay (Pdelay) is calculated periodically on each link between the slave and the MC independently. The time synchronization packet sent from the MC to all the slaves in the network is updated by each of the downstream nodes with both switching delay (the time that the packet traversed the switch) and upstream hop Pdelay.

5.2.3 PTP Domains

A domain consists of one or more PTP devices communicating with each other. PTP domain defines the scope of PTP message communication, state, operations, data sets, and timescale.

5.2.3.1 Boundary Clock

In a full E2E PTP deployment, the GMC needs to respond to each slave's delay request message. A normal profile of PTP may require a few delay calculations per second. An average GMC is capable of addressing few thousands of messages per second. This imposes that direct slave/GMC communication limits the number of overall OCs to ~8K. To scale beyond that, there is a need for a hierarchy between the GMC and the slave. This is achieved by implementing BC, either in the TOR switches or on all the switches in the DC.

The following figure shows the master/slave role that a boundary clock implements between the MC and the Slave (OC).

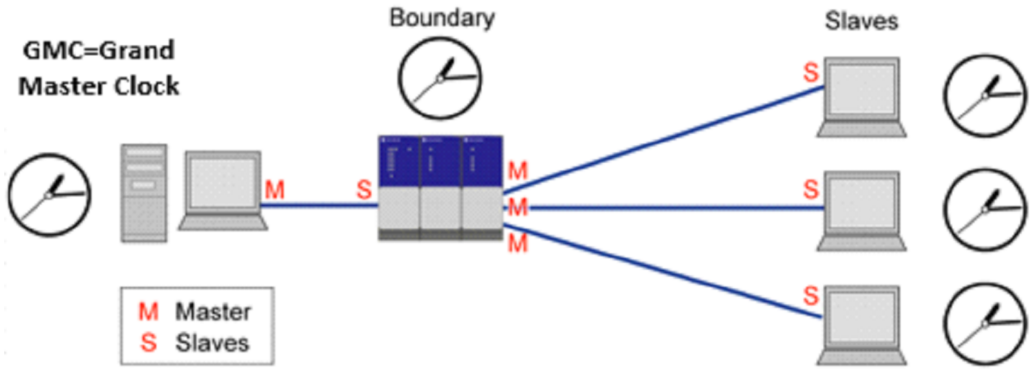

Each BC acts as a slave towards the GMC and as GMC to its local slaves. Although adding a BC device introduces accuracy degradation as explained above, it becomes mandatory when the number of slaves on a single MC exceeds few thousand devices.

Another use of BC is to bridge between networks. When running PTP over native Ethernet packets, to create larger PTP domains, there is a need to bridge between the broadcast domains. This is done by BC switches.

Default PTP Profile Attributes (SMPTE 2059-2)

Name	Range	Default
Announce interval	-3 (0.125s), 1 (2s)	-2 (0.25s)
Announce timeout interval	2, 10	3
Sync interval (logSyncInt)	-7, -1	-3
Delay request interval	logSyncInt, logSyncInt +5	logSyncInt
PTP domain	0, 127	127
Priority 1	0, 255	128
Priority 2	0, 255	128

5.2.3.2 Configuring PTP

IEEE 1588 Precision Time Protocol (PTP) may be configured either on router or switch interfaces.

To enable PTP on a router interface you could simply enable it on the selected interface.

The process of configuring PTP on a switch interface is slightly different, however. PTP should be enabled on the interface itself as well as on the respective VLAN interface(s).

All PTP configuration for switch interfaces is taken from those defined on the VLAN interface.

- Prior to enabling PTP, NTP must be disabled.
- When changing PTP configurations, PTP process restarts.
- PTP restarts under the following conditions:
    a. Removing/adding PTP-enabled VLAN from switchport configuration.
    b. Removing/adding physical interface to PTP-enabled list.
    c. Any direct PTP configuration change (forced-master, AMT, intervals, and so forth).

To configure PTP on a router interface:

1. Enable the PTP CLI commands. Run:

```
switch (config) # protocol ptp
```

2. Configure the router interface. Run:

```
switch (config) # interface ethernet 1/1 no switchport force
```

3. Add the primary IP address. Run:

```
switch (config) # interface ethernet 1/1 ip address 172.16.1.1/24
```

4. Enable PTP on the interface. Run:

```
switch (config) # interface ethernet 1/1 ptp enable
```

To verify the PTP configuration:

```
switch (config) # show ptp
PTP mode                : Boundary Clock
Message format          : Mixed
Acceptable Master Table : Enabled
Domain                  : 127
Clock identity          : 7c:fe:90:ff:fe:fa:21:88
GMC identity            : 7c:fe:90:ff:fe:fa:21:88
Number of master ports  : 1
Slave port interface    : N/A

PTP enabled interfaces:
  ------------------------------------------------------
  Port           VLAN      State        Forced Master
  ------------------------------------------------------
  Eth1/1         N/A       MASTER       no
```

To configure PTP on a switch interface:

1. Enable the PTP CLI commands. Run:

```
switch (config) # protocol ptp
```

2. Add the VLANs. Run:

```
switch (config) # vlan 2-3
```

3. Configure VLAN membership.
   For access interfaces, run:

```
switch (config) # interface ethernet 1/2 switchport mode access
switch (config) # interface ethernet 1/2 switchport access vlan 2
```

   For trunked interfaces, run:

```
switch (config) # interface ethernet 1/1 switchport mode trunk
```

4. Enable PTP on the VLAN interface. Run:

```
switch (config) # interface vlan 2 ptp enable
switch (config) # interface vlan 3 ptp enable
```

5. Enable PTP on the interface. Run:

```
switch (config) # interface ethernet 1/1 ptp enable
```

   The interface must be a member of the PTP enabled VLAN(s).

To verify the PTP configuration:

```
switch (config) # show ptp
PTP mode                : Boundary Clock
Message format          : Mixed
Acceptable Master Table : Enabled
```

```
Domain                  : 127
Clock identity          : 7c:fe:90:ff:fe:fa:21:88
GMC identity            : 7c:fe:90:ff:fe:fa:21:88
Number of master ports  : 2
Slave port interface    : N/A

PTP enabled interfaces:
  ------------------------------------------------------
  Port            VLAN      State        Forced Master
  ------------------------------------------------------
  Eth1/1          2         MASTER       no
  Eth1/2          2         MASTER       no
  Eth1/1          3         SLAVE        no
```

- Having multiple PTP connections between a pair of switches is not supported.
- Enabling PTP only for a specific VLAN on a specific trunk port is not supported.
  For example, there can be two PTP-enabled trunk ports

```
interface ethernet 1/1-1/2 switchport mode trunk
interface ethernet 1/1-1/2 switchport trunk allowed-vlan all
interface ethernet 1/1 ptp enable
interface ethernet 1/2 ptp enable
```

  and two VLAN interfaces in PTP.

```
interface vlan 10 ip address 1.1.1.1/24 primary
interface vlan 20 ip address 2.1.1.1/24 primary
interface vlan 10 ptp enable
interface vlan 20 ptp enable
```

  As the result, there will be four PTP connections:

```
switch (config) # show ptp
PTP mode                : Boundary Clock
Message format          : Mixed
Acceptable Master Table : Disabled
Domain                  : 127
TTL                     : 1
TX TS timeout (msec)    : 30
Clock identity          : 7C:FE:90:FF:FE:FF:21:08
GMC identity            : 7C:FE:90:FF:FE:FF:21:08
Number of master ports  : 2
Slave port interface    : N/A
PTP enabled interfaces:
  ------------------------------------------------------------------------------------
  Port          Po        VLAN      VRF       Transport   State        Forced Master
  ------------------------------------------------------------------------------------
  Eth1/1        N/A       10        default   IPv4        FAULTY       no
  Eth1/2        N/A       10        default   IPv4        MASTER       no
  Eth1/1        N/A       20        default   IPv4        FAULTY       no
  Eth1/2        N/A       20        default   IPv4        MASTER       no
```

  In this scenario, PTP cannot be enabled on VLAN 10 for port 1/1 and VLAN 20 for port 1/2.
  If PTP is disabled on VLAN 10 or 20, it will disable PTP for both ports 1/1 and 1/2.

5.2.4 Securing PTP Infrastructure

To protect the switch from rogue or mis-configured PTP endpoints, you may secure your Boundary Clock ports by creating an Acceptable Master Table (AMT) and configuring known PTP ports to always behave as a master port via the Forced Master option.

The AMT is a whitelist of up to 8 clock identities that are admissible to take part as valid GrandMasters in the Best Master Clock Algorithm (BMCA).

The Forced Master is enabled on a per-port basis to prevent processing announce messages from a PTP endpoint connected to it, in order for it to always stay in a Master state.

To configure Forced Master on a switch interface, you must enable it on the interface itself as well as on the respective VLAN interface(s).

To configure Acceptable Master Table, add the validated clock identities:

```
switch (config) # ptp amt E4:1D:2D:FF:FE:46:13:88
switch (config) # ptp amt E4:1D:2D:FF:FE:44:23:B7
```

To verify the Acceptable Master Table configuration:

```
switch (config) # show ptp amt

Clock Identities:
  E4:1D:2D:FF:FE:44:23:B7
  E4:1D:2D:FF:FE:46:13:88
```

To enable Forced Master on a router interface:

```
switch (config) # interface ethernet 1/2 ptp enable forced-master
```

To verify PTP configuration:

```
switch (config) # show ptp
PTP mode                 : Boundary Clock
Message format           : Mixed
Acceptable Master Table  : Enabled
Domain                   : 127
Clock identity           : 7c:fe:90:ff:fe:fa:21:88
GMC identity             : 7c:fe:90:ff:fe:fa:21:88
Number of master ports   : 1
Slave port interface     : N/A

PTP enabled interfaces:
  ------------------------------------------------------
  Port           VLAN      State        Forced Master
  ------------------------------------------------------
  Eth1/2         N/A       MASTER       yes
```

To configure Forced Master on a switch interface:

1. Enable Forced Master on the VLAN interface. Run:

   ```
   switch (config) # interface vlan 2 ptp enable forced-master
   ```

2. Enable Forced Master on the interface. Run:

   ```
   switch (config) # interface ethernet 1/1 ptp enable forced-master
   ```

   The interface should be a member in the PTP enabled VLAN(s).

To verify PTP configuration:

```
switch (config) # show ptp
PTP mode                 : Boundary Clock
Message format           : Mixed
Acceptable Master Table  : Enabled
Domain                   : 127
Clock identity           : 7c:fe:90:ff:fe:fa:21:88
GMC identity             : 7c:fe:90:ff:fe:fa:21:88
Number of master ports   : 2
Slave port interface     : N/A

PTP enabled interfaces:
  ------------------------------------------------------
  Port           VLAN      State        Forced Master
  ------------------------------------------------------
  Eth1/1         2         MASTER       yes
  Eth1/1         3         SLAVE        no
```

Forced Master is indicated as “yes” only if enabled on the interface and the corresponding VLAN interface.

5.2.5 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following:

- IEEE 1588 Precision Time Protocol Design Guide → Guides (scroll to the bottom of the page)
- IEEE 1588 PTP on Spectrum Switches

5.2.6 PTP Commands

- 5.2.6.1 protocol ptp
- 5.2.6.2 ptp amt
- 5.2.6.3 ptp announce interval
- 5.2.6.4 ptp announce timeout
- 5.2.6.5 ptp delay-req interval
- 5.2.6.6 ptp domain
- 5.2.6.7 ptp enable
- 5.2.6.8 ptp enable forced-master
- 5.2.6.9 ptp enable forced-master allow-mgmt
- 5.2.6.10 ptp enable ipv6
- 5.2.6.11 ptp mean-path-delay
- 5.2.6.12 ptp message-format
- 5.2.6.13 ptp monitor interval
- 5.2.6.14 ptp monitor interval phc
- 5.2.6.15 ptp monitor logging enable
- 5.2.6.16 ptp offset-from-master
- 5.2.6.17 ptp priority
- 5.2.6.18 ptp sync interval
- 5.2.6.19 ptp tll
- 5.2.6.20 clear ptp amt log
- 5.2.6.21 clear ptp forced-master log
- 5.2.6.22 clear ptp interface counters
- 5.2.6.23 clear ptp timeout counters
- 5.2.6.24 clear ptp vrf counters
- 5.2.6.25 ptp vrf enable
- 5.2.6.26 show ptp
- 5.2.6.27 show ptp monitor
- 5.2.6.28 show ptp monitor phc
- 5.2.6.29 show ptp timeout counters
- 5.2.6.30 show ptp vrf
- 5.2.6.31 show ptp vrf counters
- 5.2.6.32 show ptp amt

- 5.2.6.33 show ptp interface port-channel
- 5.2.6.34 show ptp interface port-channel counters
- 5.2.6.35 show ptp amt log
- 5.2.6.36 show ptp clock
- 5.2.6.37 show ptp clock parent
- 5.2.6.38 show ptp forced-master
- 5.2.6.39 show ptp
- 5.2.6.40 show ptp clock foreign-masters
- 5.2.6.41 show ptp interface ethernet counters
- 5.2.6.42 show ptp interface
- 5.2.6.43 show ptp interface ethernet
- 5.2.6.44 show ptp interface vlan
- 5.2.6.45 show ptp interface vlan ethernet
- 5.2.6.46 show ptp interface vlan counters
- 5.2.6.47 show ptp interface vlan ethernet counters
- 5.2.6.48 show ptp time-property
- 5.2.6.49 show ptp status
- 5.2.6.50 PTP Debuggability Logging Examples
  - 5.2.6.50.1 Change of the State of Particular PTP Port
  - 5.2.6.50.2 Change of Grandmaster Clock

5.2.6.1 protocol ptp

	protocol ptp Enables PTP on the switch.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.4110

Example	`switch (config) # protocol ptp`
	`...`
	`switch (config) #`

Related Commands	
Notes	

5.2.6.2 ptp amt

	ptp amt <clock-id> no ptp amt <clock-id> Adds an acceptable master table entry. The no form of the command removes an acceptable master entry.	
Syntax Description	clock-id	Clock ID
Default	N/A	
Configuration Mode	config	
History	3.6.8100	

Example	`switch (config) # ptp amt 00:11:22:FF:FE:33:44:55:66`
Related Commands	show ptp amt show ptp amt log show ptp clock
Notes	

5.2.6.3 ptp announce interval

	ptp announce interval <interval> Configures PTP announce interval.	
Syntax Description	interval	Range: -3 to 1 Default: -2
Default	N/A	
Configuration Mode	config interface port-channel config interface ethernet config interface vlan	
History	3.6.4110	
	3.6.8008	Added “interface vlan” configuration mode
	3.6.8100	Added "interface port-channel" configuration mode

Example	`switch (config 1/1) # ptp announce interval -2`
	`...`
	`switch (config 1/1) #`	

Related Commands	show ptp interface show ptp interface <ethernet | port-channel | vlan>	
Notes		

5.2.6.4 ptp announce timeout

	ptp announce timeout <timeout> Configures PTP announce timeout.	
Syntax Description	timeout	Range: 2-10 Default: 3
Default	N/A	
Configuration Mode	config interface port-channel config interface ethernet config interface vlan	
History	3.6.4110	
	3.6.8008	Added “interface vlan” configuration mode
	3.6.8100	Added "interface port-channel" configuration mode

Example	`switch (config 1/1) # ptp announce timeout 3`
	`...`
	`switch (config 1/1) #`	

Related Commands	show ptp interface show ptp interface <ethernet | port-channel | vlan>	

Notes	

5.2.6.5 ptp delay-req interval

	ptp delay-req interval <interval> Configures PTP delay-req interval.	
Syntax Description	interval	Range: 0-5 Default: 0
Default	N/A	
Configuration Mode	config interface port-channel config interface ethernet config interface vlan	
History	3.6.4110	
	3.6.8008	Added “interface vlan” configuration mode
	3.8.8100	"interface port-channel" configuration mode
	3.9.0600	Updated example and added note

Example	`switch (config 1/1) # ptp delay-req interval 0`
	`...`
	`switch (config 1/1) #`	

Related Commands	show ptp interface show ptp interface <ethernet | port-channel | vlan>	
Notes	IEEE 1588 defines delay-request as an offset from Sync Interval (logSyncInt). A value of 0 therefore matches the defined logSyncInt value. Example: logSyncInt = -3, delay-req = 0 implies delay-req message rate is -3	

5.2.6.6 ptp domain

	ptp domain <domain number> Inserts the number of ptp domain.	
Syntax Description	domain number	Range: 0-127
Default	127	
Configuration Mode	config	
History	3.6.4110	

Example	`switch (config) # ptp domain`
	`...`
	`switch (config) #`	

Related Commands	show ptp clock	
Notes		

5.2.6.7 ptp enable

	ptp enable no ptp enable Enables PTP per interface. The no form of the command disables PTP per interface.	
Syntax Description	N/A	
Default	no ptp enable	
Configuration Mode	config interface ethernet config interface port-channel config interface vlan	
History	3.6.4110	
	3.6.8008	Added “config interface vlan” configuration mode
	3.6.8100	Added “config interface port-channel” configuration mode

Example	`switch (config interface ethernet 1/1) # ptp enable`
	`...`
	`switch (config interface ethernet 1/1) #`	

Related Commands	show ptp show ptp interface show ptp interface <ethernet | port-channel | vlan>	
Notes		

5.2.6.8 ptp enable forced-master

	ptp enable forced-master no ptp enable forced-master Configures PTP interfaces to forced master state. The command allows dropping both announce and mgmt PTP messages from other PTP peers in the network. The no form of the command removes PTP interfaces from forced master state.	
Syntax Description	N/A	
Default	no ptp enable forced-master	
Configuration Mode	config interface ethernet config interface vlan config interface port-channel	
History	3.6.8100	
	3.9.2300	Added note and modified command description.
Example	`switch (config interface ethernet 1/1) # ptp enable forced-master`	
Related Commands	show ptp show ptp interface show ptp interface <ethernet | port-channel | vlan> ptp enable forced-master allow-mgmt	
Notes	In order to enable forced-master on the VLAN/LAG, it should be enabled on the VLAN/LAG members as well as on the VLAN/LAG itself.	

5.2.6.9 ptp enable forced-master allow-mgmt

	ptp enable forced-master allow-mgmt no ptp enable forced-master allow-mgmt Configures PTP interfaces to forced master state. The command drops only announce PTP messages and allows passing mgmt PTP messages from other PTP peers in the network. The no form of the command removes PTP interfaces from forced master state.
Syntax Description	N/A
Default	no ptp enable forced-master allow-mgmt
Configuration Mode	config interface ethernet config interface vlan config interface port-channel
History	3.9.2300
Example	`switch (config 1/1) # ptp enable forced-master allow-mgmt`
Related Commands	ptp enable forced-master
Notes	In order to enable forced-master allow-mgmt on the VLAN/LAG, it should be enabled on the VLAN/LAG members as well as on the VLAN/LAG itself.

5.2.6.10 ptp enable ipv6

	ptp enable [forced-master] [ipv6 [mcast-scope link-local]] no ptp enable [forced-master] [ipv6 [mcast-scope link-local]] Configures PTP on the ethernet interface and enables the forced-master and support of IPv6 with a specified scope. The no form of the command removes the support from the interface.	
Syntax Description	mcast-scope link-local	Sets the IPv6 multicast scope to link-local.
Default	no ptp enable ipv6	
Configuration Mode	config interface ethernet	
History	3.8.2000	
Example	`switch (config interface ethernet 1/1) # ptp enable ipv6 mcast-scope link-local`	
Related Commands	show ptp	
Notes	When configuring PTP IPv6, the "global" multicast scope is the default.	

5.2.6.11 ptp mean-path-delay

	ptp mean-path-delay <value> no ptp mean-path-delay <value> Enables logging of the mean path delay value if it exceeds the specified threshold. Disables logging of the mean path delay value if it exceeds the specified threshold.	
Syntax Description	value	10-1000000000 (ns). Default 1000000000
Default	Enabled	
Configuration Mode	config	

History	3.8.2100
Example	`switch (config) # ptp mean-path-delay 10000000`
Logging Examples	Example of ptp mean-path-delay 10: Nov 11 16:18:04 arc-switch142 ptp4l: [3083.530] PTP [Debuggability]: PTP Grandmaster clock has changed from ec0d9a.fffe.603848 to 248a07.fffe.9e9adc Nov 11 16:18:04 arc-switch142 ptp4l: [3083.530] port 1: Interface Eth1/10 state changed from MASTER to UNCALIBRATED on RS_SLAVE Nov 11 16:18:05 arc-switch142 ptp4l: [3084.404] PTP slave port Eth1/10 High offset from Master -58705983752 (ns) Nov 11 16:18:06 arc-switch142 ptp4l: [3084.904] PTP slave port Eth1/10 High offset from Master -58705990066 (ns) Nov 11 16:18:06 arc-switch142 ptp4l: [3085.062] PTP slave port Eth1/10 High Mean Path Delay 56 (ns) Nov 11 16:18:06 arc-switch142 ptp4l: [3085.225] PTP slave port Eth1/10 High Mean Path Delay 313 (ns) Nov 11 16:18:06 arc-switch142 ptp4l: [3085.318] PTP slave port Eth1/10 High Mean Path Delay 709 (ns) Nov 11 16:18:06 arc-switch142 ptp4l: [3085.404] PTP slave port Eth1/10 High offset from Master -58705997158 (ns) Nov 11 16:18:07 arc-switch142 ptp4l: [3085.904] port 1: Interface Eth1/10 state changed from UNCALIBRATED to SLAVE on MASTER_CLOCK_SELECTED Nov 11 16:18:07 arc-switch142 ptp4l: [3085.966] PTP slave port Eth1/10 High Mean Path Delay 709 (ns) Nov 11 16:18:07 arc-switch142 ptp4l: [3086.192] PTP slave port Eth1/10 High Mean Path Delay 709 (ns) Nov 11 16:18:07 arc-switch142 ptp4l: [3086.215] PTP slave port Eth1/10 High Mean Path Delay 709 (ns) Nov 11 16:18:07 arc-switch142 ptp4l: [3086.240] PTP slave port Eth1/10 High Mean Path Delay 709 (ns) Nov 11 16:18:07 arc-switch142 ptp4l: [3086.244] PTP slave port Eth1/10 High Mean Path Delay 246 (ns) Nov 11 16:18:07 arc-switch142 ptp4l: [3086.404] port 1: Interface Eth1/10 state changed from SLAVE to UNCALIBRATED on SYNCHRONIZATION_FAULT Nov 11 16:18:09 arc-switch142 ptp4l: [3087.904] port 1: Interface Eth1/10 state changed from UNCALIBRATED to SLAVE on MASTER_CLOCK_SELECTED Nov 11 16:19:10 arc-switch142 ptp4l: [3090.711] PTP slave port Eth1/10 High Mean Path Delay 15 (ns) Nov 11 16:19:10 arc-switch142 ptp4l: [3090.740] PTP slave port Eth1/10 High Mean Path Delay 15 (ns) Nov 11 16:19:10 arc-switch142 ptp4l: [3090.831] PTP slave port Eth1/10 High Mean Path Delay 23 (ns) Nov 11 16:19:10 arc-switch142 ptp4l: [3090.879] PTP slave port Eth1/10 High Mean Path Delay 23 (ns) Nov 11 16:19:10 arc-switch142 ptp4l: [3091.025] PTP slave port Eth1/10 High Mean Path Delay 23 (ns) Nov 11 16:19:11 arc-switch142 ptp4l: [3091.128] PTP slave port Eth1/10 High Mean Path Delay 21 (ns) Nov 11 16:19:11 arc-switch142 ptp4l: [3091.292] PTP slave port Eth1/10 High Mean Path Delay 20 (ns) Nov 11 16:19:11 arc-switch142 ptp4l: [3091.406] PTP slave port Eth1/10 High Mean Path Delay 20 (ns) Nov 11 16:19:11 arc-switch142 ptp4l: [3091.621] PTP slave port Eth1/10 High Mean Path Delay 20 (ns) Nov 11 16:19:11 arc-switch142 ptp4l: [3091.625] PTP slave port Eth1/10 High Mean Path Delay 20 (ns)
Related Commands	show ptp clock show ptp status show log

Notes	If the mean path delay exceeds the threshold, the following ptp4l log message will appear: “Oct 11 19:04:41 arc-switch142 ptp4l: [242.721] PTP slave port Eth1/10 High Mean Path Delay 65536 (ns)”

5.2.6.12 ptp message-format

	ptp message-format {mixed | multicast} Configures PTP delay request messages format.	
Syntax Description	mixed	Sends unicast delay request packets
	multicast	Sends multicast delay request packets
Default	mixed	
Configuration Mode	config	
History	3.6.8008	
Example	`switch (config) # ptp message-format mixed`	
Related Commands		
Notes		

5.2.6.13 ptp monitor interval

	ptp monitor interval <interval> no ptp monitor interval Configure the time interval in which summary statistics of the clock are printed. It is specified as a power of two in seconds. The no form of the command sets monitor interval to its default value.	
Syntax Description	interval	Range: 0 to 20, power of two in seconds. For example, when configuring the interval to 3, the time interval will be configured to 8 seconds (2^3= 8 seconds). Default: 0 (2^0=1 second)
Default	ptp monitor interval 0	
Configuration Mode	config	
History	3.10.2000	

Example	switch (config) # ptp monitor interval 1 switch (config) # show ptp monitor PTP monitor logging : disabled PTP monitor interval (seconds) : 2 PTP monitor PHC interval (seconds) : 1.0 --- --- Interface Time RMS max freq mean freq delay mean delay --- --- Eth1/13 2022/03/02 18:37:22.906 2 2 +33 2 109 0 Eth1/13 2022/03/02 18:37:20.906 4 7 +37 4 108 1 switch (config) # ptp monitor interval 3 switch (config) # show ptp monitor PTP monitor logging : disabled PTP monitor interval (seconds) : 8 PTP monitor PHC interval (seconds) : 1.0 --- --- Interface Time RMS max freq mean freq delay mean delay --- --- Eth1/13 2022/03/02 18:37:28.305 1 2 +23 3 104 0 Eth1/13 2022/03/02 18:37:20.306 3 7 +22 4 103 1
Related Commands	ptp monitor interval phc ptp monitor logging enable show ptp monitor show ptp monitor phc
Notes	Configure the time interval in which summary statistics of the clock are printed both in the ‘show ptp monitor’ output and into the debug log from ptp4l process (if ‘ptp monitor logging enable’ is configured).

5.2.6.14 ptp monitor interval phc

	ptp monitor interval phc <interval> no ptp monitor interval phc Configure the time interval in which summary statistics of the PHC clock are printed. It is specified in seconds. The no form of the command sets PHC monitor interval to its default value.	
Syntax Description	interval	Range: 0.0 to 100000.0 seconds Default: 1.0 second
Default	ptp monitor interval phc 1.0	
Configuration Mode	config	
History	3.10.2000	

Example	switch (config) # ptp monitor interval phc 10 switch (config) # show ptp monitor phc PTP monitor logging : disabled PTP monitor interval (seconds) : 2 PTP monitor PHC interval (seconds) : 10.0 --- Interface Time RMS max freq mean freq delay mean delay --- Eth1/13 2022/03/02 18:37:22.506 5 5 +33 3 111 0 Eth1/13 2022/03/02 18:37:12.506 8 8 +32 5 118 1 switch (config) # ptp monitor interval phc 0.5 switch (config) # show ptp monitor phc PTP monitor logging : disabled PTP monitor interval (seconds) : 1 PTP monitor PHC interval (seconds) : 0.5 --- Interface Time RMS max freq mean freq delay mean delay --- Eth1/13 2022/03/02 18:37:22.906 4 4 +13 4 101 1 Eth1/13 2022/03/02 18:37:22.406 7 8 +22 2 113 0
Related Commands	ptp monitor interval ptp monitor logging enable show ptp monitor show ptp monitor phc
Notes	Configure the time interval in which summary statistics of the clock are printed both in the ‘show ptp monitor phc’ output and into the debug log from phc2sys process (if ‘ptp monitor logging enable’ is configured).

5.2.6.15 ptp monitor logging enable

	ptp monitor logging enable no ptp monitor logging enable Enables PTP monitor logging into ‘/var/log/debug’ file. The no form of the command disables PTP monitor logging into ‘/var/log/debug’ file.
Syntax Description	N/A
Default	Disabled By default, 1 ptp4l and phc2sys message with clock statistics is printed per 1 second.
Configuration Mode	config interface ethernet
History	3.10.2000
Example	switch (config) # ptp monitor logging enable switch (config) # show log debug matching rms [ptp4l.INFO]: [4628.609] rms 11 max 34 freq +37 +/- 13 delay 109 +/- 1 [phc2sys.INFO]: [4628.974] CLOCK_REALTIME rms 379 max 1651 freq -16541 +/- 684 delay 1391+/- 96

Related Commands	ptp monitor interval ptp monitor interval phc show ptp monitor show ptp monitor phc
Notes	PTP monitor logs include such parameters for ptp4l and phc2sys clock: RMS: offset root mean square Max: maximum absolute offset Freq mean offset: frequency mean deviation Freq offset: frequency standard deviation Path delay mean: mean patch delay deviation Path delay: standard path delay deviation

5.2.6.16 ptp offset-from-master

	ptp offset-from-master <value> <value> Enables logging of the offset from master value if it exceeds the specified threshold.	
Syntax Description	values	[-1000000000; -10] [10; 1000000000]. Default [-100000; -10] [10; 100000]
Default	Enabled	
Configuration Mode	config	
History	3.8.2100	
Example	`switch (config) # ptp offset-from-master -100 2345`	

Logging Example	Example of ptp offset-from-master -10 10: Nov 11 16:09:54 arc-switch142 ptp4l: [2593.020] port 1: Interface Eth1/10 state changed from MASTER to UNCALIBRATED on RS_SLAVE Nov 11 16:09:54 arc-switch142 ptp4l: [2593.269] port 1: Interface Eth1/10 state changed from UNCALIBRATED to SLAVE on MASTER_CLOCK_SELECTED Nov 11 16:10:03 arc-switch142 ptp4l: [2601.897] PTP slave port Eth1/10 High offset from Master -11 (ns) Nov 11 16:10:03 arc-switch142 ptp4l: [2602.022] PTP slave port Eth1/10 High offset from Master -14 (ns) Nov 11 16:10:03 arc-switch142 ptp4l: [2602.272] PTP slave port Eth1/10 High offset from Master -11 (ns) Nov 11 16:10:03 arc-switch142 ptp4l: [2602.397] PTP slave port Eth1/10 High offset from Master -13 (ns) Nov 11 16:10:14 arc-switch142 ptp4l: [2613.526] PTP slave port Eth1/10 High offset from Master -11 (ns) Nov 11 16:10:21 arc-switch142 ptp4l: [2620.279] PTP slave port Eth1/10 High offset from Master 12 (ns) Nov 11 16:10:21 arc-switch142 ptp4l: [2620.529] PTP slave port Eth1/10 High offset from Master 12 (ns) Nov 11 16:10:28 arc-switch142 ptp4l: [2627.656] PTP slave port Eth1/10 High offset from Master -11 (ns) Nov 11 16:10:29 arc-switch142 ptp4l: [2627.907] PTP slave port Eth1/10 High offset from Master -11 (ns) Nov 11 16:10:52 arc-switch142 ptp4l: [2650.790] PTP slave port Eth1/10 High offset from Master -13 (ns) Nov 11 16:11:01 arc-switch142 ptp4l: [2660.419] PTP slave port Eth1/10 High offset from Master 11 (ns) Nov 11 16:11:13 arc-switch142 ptp4l: [2672.548] PTP slave port Eth1/10 High offset from Master -13 (ns) Nov 11 16:11:17 arc-switch142 ptp4l: [2676.674] PTP slave port Eth1/10 High offset from Master 11 (ns) Nov 11 16:11:21 arc-switch142 ptp4l: [2680.676] PTP slave port Eth1/10 High offset from Master 11 (ns) Nov 11 16:11:24 arc-switch142 ptp4l: [2683.552] PTP slave port Eth1/10 High offset from Master -11 (ns) Nov 11 16:11:28 arc-switch142 ptp4l: [2687.553] PTP slave port Eth1/10 High offset from Master -11 (ns) Nov 11 16:11:34 arc-switch142 ptp4l: [2692.930] PTP slave port Eth1/10 High offset from Master -11 (ns) Nov 11 16:11:44 arc-switch142 ptp4l: [2703.059] PTP slave port Eth1/10 High offset from Master 12 (ns) Nov 11 16:11:44 arc-switch142 ptp4l: [2703.309] PTP slave port Eth1/10 High offset from Master 11 (ns) Nov 11 16:11:50 arc-switch142 ptp4l: [2709.561] PTP slave port Eth1/10 High offset from Master -11 (ns) Nov 11 16:11:55 arc-switch142 ptp4l: [2713.937] PTP slave port Eth1/10 High offset from Master -13 (ns) Nov 11 16:11:55 arc-switch142 ptp4l: [2714.062] PTP slave port Eth1/10 High offset from Master -15 (ns) Nov 11 16:11:55 arc-switch142 ptp4l: [2714.312] PTP slave port Eth1/10 High offset from Master -14 (ns) Nov 11 16:11:55 arc-switch142 ptp4l: [2714.438] PTP slave port Eth1/10 High offset from Master -11 (ns)
Related Commands	show log show ptp clock show ptp status
Notes	If the mean path delay exceeds the threshold, the following ptp4l log message will appear: “Oct 11 19:04:41 arc-switch142 ptp4l: [242.721] PTP slave port Eth1/10 High offset from Master 36766720739 (ns)”

5.2.6.17 ptp priority

	ptp priority{1 | 2} <priority> Configures PTP primary priority.	
Syntax Description	priority	Range: 0-255
Default	128	

Configuration Mode	config
History	3.6.4110

Example	`switch (config) # ptp priority1 128`
	`...`
	`switch (config) #`

Related Commands	show ptp clock
Notes	

5.2.6.18 ptp sync interval

	ptp sync interval <interval> Configures PTP sync interval.	
Syntax Description	interval	Range: -7 to -1 Default: -3
Default	N/A	
Configuration Mode	config interface port-channel config interface ethernet config interface vlan	
History	3.6.4110	
	3.6.8008	Added “interface vlan” configuration mode
	3.6.8100	Added "interface port-channel" configuration mode

Example	`switch (config 1/1) # ptp sync interval -3`
	`...`
	`switch (config 1/1) #`	

Related Commands	show ptp interface show ptp interface <ethernet | port-channel | vlan>	
Notes		

5.2.6.19 ptp ttl

	ptp ttl <ttl_value> no ptp ttl Sets the TTL value of the PTP messages. The no form of the command sets the PTP UDP TTL value back to its default value of 1.	
Syntax Description	ttl_value	1-255
Default	PTP TTL is 1 by default.	
Configuration Mode	config	
History	3.9.2000	

Example	`switch (config) # ptp ttl 10`
	`switch (config) # show ptp`
	`PTP mode : Boundary Clock`
	`Message format : Mixed`
	`Acceptable Master Table : Disabled`
	`Domain : 127`
	`TTL : 10`	

Related Commands	show ptp
Notes	

5.2.6.20 clear ptp amt log

	clear ptp amt log Clears log of received clock IDs outside of acceptable master table.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.8100
Example	`switch (config) # clear ptp amt log`
Related Commands	show ptp amt show ptp amt log
Notes	

5.2.6.21 clear ptp forced-master log

	clear ptp forced-master log Clears log of received clock IDs on forced master interface.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.8100
Example	`switch (config) # clear ptp forced-master log`
Related Commands	show ptp forced-master show ptp forced-master log
Notes	

5.2.6.22 clear ptp interface counters

	clear ptp interface [vlan <id>] [port-channel <id>] [ethernet <slot>/<port>[/<subport>]] counters Clears PTP counters for specified VLAN member interface.	
Syntax Description	id	VLAN or LAG ID
	<slot>/<port>/<subport>	Ethernet port ID (e.g. 1/3/1)
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	
	3.8.2000	Added example

Example	`switch (config 1/1) # clear ptp interface vlan 2 ethernet 1/1 counters`
Related Commands	show ptp interface <ethernet | port-channel | vlan> counters
Notes	

5.2.6.23 clear ptp timeout counters

	clear ptp timeout counters Clears global PTP timeout counters.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.10.2000
Example	`switch (config) # clear ptp timeout counters`
Related Commands	show ptp timeout counters
Notes	

5.2.6.24 clear ptp vrf counters

	clear ptp vrf <vrf-name> counters Clears the PTP VRF counters.	
Syntax Description	vrf-name	Name of PTP enabled VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.7.1000	
Example	`switch (config) # clear ptp vrf cust1 counters`	
Related Commands	show ptp vrf counters	
Notes	This command clears interface statistics on all PTP enabled interfaces in a specific PTP enabled VRF.	

5.2.6.25 ptp vrf enable

	ptp vrf <vrf-name> enable [forced-master] no ptp vrf <vrf-name> enable [forced-master] This command enables PTP in VRF. Running the no form of this command disables PTP in a specified VRF.
Syntax Description	N/A
Default	N/A
Configuration Mode	Configure terminal
History	3.7.1000
Example	*switch (config) # ptp vrf cust1 enable forced-master*

Related Commands	show ptp show ptp vrf show ptp forced-master show ptp vrf counters clear ptp vrf counters ptp vrf announce interval ptp vrf announce timeout ptp vrf delay-req interval ptp vrf sync interval
Related Commands	PTP needs to be enabled on interfaces in VRF as well.

5.2.6.26 show ptp

	show ptp Displays PTP configuration and operation data.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4110	
	3.6.8008	Updated example
	3.6.8100	Updated example
	3.8.2000	Updated example
	3.9.2000	Updated example, adding TTL field

Example	`switch (config) # show ptp`
	`PTP mode : Boundary Clock`
	`Message format : Mixed`
	`Acceptable Master Table : Disabled`
	`Domain : 127`
	`TTL : 10`
	`Clock identity : 7c:fe:90:ff:fe:fa:23:88`
	`GMC identity : 7c:fe:90:ff:fe:fa:23:88`
	`Number of master ports : 0`
	`Slave port interface : N/A`
	
	`PTP enabled interfaces:`
	`Port Po VLAN VRF Transport State Forced Master`
	`Eth1/1 N/A N/A default IPv4 SLAVE no`
	`Eth1/2 N/A N/A default IPv6 MASTER no`	

Related Commands		
Notes		

5.2.6.27 show ptp monitor

	show ptp monitor Displays last 100 entries of the PTP clock monitor data from ptp4l process.
Syntax Description	N/A
Default	By default, 1 entry with clock monitor data is printed per 1 second.
Configuration Mode	Any command mode

History	3.10.2000
Example	switch (config) # show ptp monitor PTP monitor logging : disabled PTP monitor interval (seconds) : 1 PTP monitor PHC interval (seconds): 1.0 ---------- ---------- Interface Time RMS max freq mean freq delay mean delay ---------- ---------- Eth1/13 2022/03/02 18:37:20.906 3 4 +31 4 110 1 Eth1/13 2022/03/02 18:37:19.906 4 8 +29 5 111 0 Eth1/13 2022/03/02 18:37:18.906 3 6 +33 4 109 0 Eth1/13 2022/03/02 18:37:17.905 4 7 +32 4 109 0 Eth1/13 2022/03/02 18:37:16.905 3 4 +36 4 109 0
Related Commands	ptp monitor logging enable ptp monitor interval ptp monitor interval phc show ptp monitor phc
Notes	PTP monitor logs include such clock parameters from ptp4l process: RMS: offset root mean square Max: maximum absolute offset Freq mean offset: frequency mean deviation Freq offset: frequency standard deviation Path delay mean: mean patch delay deviation Path delay: standard path delay deviation

5.2.6.28 show ptp monitor phc

	show ptp monitor phc Displays last 100 entries of the PTP clock monitor data from phc2sys process.
Syntax Description	N/A
Default	By default, 1 entry with clock monitor data is printed per 1 second.
Configuration Mode	Any command mode
History	3.10.2000

Example	switch (config) # show ptp monitor phc PTP monitor logging : disabled PTP monitor interval (seconds) : 1 PTP monitor PHC interval (seconds): 1.0 --- --- Interface Time RMS max freq mean freq delay mean delay --- --- Eth1/13 2022/03/02 18:37:09.906 2 2 +33 2 109 0 Eth1/13 2022/03/02 18:37:08.906 4 7 +37 4 108 1 Eth1/13 2022/03/02 18:37:07.906 3 4 +21 4 110 1 Eth1/13 2022/03/02 18:37:06.905 3 7 +32 4 109 0 Eth1/13 2022/03/02 18:37:05.905 3 4 +36 4 109 0
Related Commands	ptp monitor logging enable ptp monitor interval ptp monitor interval phc show ptp monitor
Notes	PTP monitor logs include such clock parameters from phc2sys process: RMS: offset root mean square Max: maximum absolute offset Freq mean offset: frequency mean deviation Freq offset: frequency standard deviation Path delay mean: mean patch delay deviation Path delay: standard path delay deviation

5.2.6.29 show ptp timeout counters

	show ptp timeout counters Displays global PTP timeout counters.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.10.2000
Example	switch (config) # show ptp timeout counters Rx: 0 announce timeout count 0 sync timeout count Tx: 0 send announce failed count 0 send sync failed count 0 send follow up failed count 0 send delay response failed count 0 send delay request failed count
Related Commands	clear ptp timeout counters
Notes	• RX timeout counters are increased when we fail to receive announce or sync PTP message in the configured interval. • TX timeout (send failed) counters are increased when we fail to send specific PTP message.

5.2.6.30 show ptp vrf

	show ptp vrf <vrf_name> Displays interfaces in VRF PTP related data.	
Syntax Description	vrf-name	Name of PTP enabled VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.7.1000	
	3.8.2000	Updated example

Example	`switch (config) # show ptp vrf`
	`Interface name: Eth1/1`
	`Channel group ID: N/A`
	`VRF: cust1`
	`IP Address: 1.1.1.1`
	`Port Clock identity: E4:1D:2D:FF:FE:44:65:C8`
	`PTP Port number: 1`
	`PTP operational state: UP`
	`PTP interface state: MASTER`
	`Forced Master: no`
	`Delay request interval(log mean): 0`
	`Announce receipt time out: 3`
	`Announce interval(log mean): -2`
	`Sync interval(log mean): -3`
	`Delay Mechanism: End to End`
	`Transport protocol: UDP IPv4`
	`IPv6 Multicast scope ID: N/A`
	
	`Interface name: Eth1/2`
	`Channel group ID: N/A`
	`VRF: default`
	`IP Address: 2.2.2.2`
	`Port Clock identity: E4:1D:2D:FF:FE:44:65:C8`
	`PTP Port number: 1`
	`PTP interface state: SLAVE`
	`PTP operational state: UP`
	`Forced Master: no`
	`Delay request interval(log mean): 0`
	`Announce receipt time out: 3`
	`Announce interval(log mean): -2`
	`Sync interval(log mean): -3`
	`Delay Mechanism: End to End`
	`Transport protocol: UDP IPv4`
	`IPv6 Multicast scope ID: N/A`
	
	`Interface name: Eth1/1`
	`Channel group ID: N/A`
	`VRF: cust1`
	`IP Address: 1.1.1.1`
	`Port Clock identity E4:1D:2D:FF:FE:44:65:C8`
	`PTP Port number: 1`
	`PTP interface state: MASTER`
	`Forced Master: no`
	`Delay request interval(log mean): 0`
	`Announce receipt time out: 3`
	`Announce interval(log mean): -2`
	`Sync interval(log mean): -3`
	`Delay Mechanism: End to End`
	`Transport protocol: UDP IPv4`
	`IPv6 Multicast scope ID: N/A`	

Related Commands		
Notes	Displays ptp state of all PTP-enabled interfaces in all PTP-enabled VRFs.	

5.2.6.31 show ptp vrf counters

	show ptp vrf <vrf-name> counters Displays port statistics on interfaces in VRF.	
Syntax Description	vrf-name	Name of PTP enabled VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.7.1000	
Example	switch (config) # show ptp vrf cust1 counters VRF: cust1 Eth1/1 RX 0 Sync message count 0 Delay request message count 0 PDelay request message count 0 PDelay response message count 0 Follow Up message count 0 Delay response message count 0 PDelay response follow Up message count 0 Announce message count 0 Signalling message count 0 Management message count TX 0 Sync message count 0 Delay request message count 0 PDelay request message count 0 PDelay response message count 0 Follow Up message count 0 Delay response message count 0 PDelay response follow Up message count 0 Announce message count 0 Signalling message count 0 Management message count 0 Forwarded Management message count Eth1/2 RX 0 Sync message count 0 Delay request message count 0 PDelay request message count 0 PDelay response message count 0 Follow Up message count 0 Delay response message count 0 PDelay response follow Up message count 0 Announce message count 0 Signalling message count 0 Management message count TX 0 Sync message count 0 Delay request message count 0 PDelay request message count 0 PDelay response message count 0 Follow Up message count 0 Delay response message count 0 PDelay response follow Up message count 0 Announce message count 0 Signalling message count 0 Management message count 0 Forwarded Management message count	
Related Commands		
Notes	Display ptp counters of all PTP enabled interfaces in specific PTP enabled VRF.	

5.2.6.32 show ptp amt

	show ptp amt Displays acceptable master table.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.8100

Example	`switch (config) # show ptp amt`
	`Clock Identities:`
	`00:11:22:ff:fe:44:55:66`
	`66:55:44:ff:fe:22:11:00`

Related Commands	show ptp amt log clear ptp amt log
Notes	

5.2.6.33 show ptp interface port-channel

	show ptp interface port-channel <po-id> Displays LAG member interfaces PTP related data.	
Syntax Description	po-id	LAG ID
Default	N/A	
Configuration Mode	Any command mode	
History	3.7.1000	
	3.8.2000	Updated example

Example	switch (config) # show ptp interface port-channel 3 Interface name: Eth1/10 Channel group ID: 3 VRF: default IP Address: 1111:0:0:0:0:0:0:0/64 Port Clock identity: ec:0D:9a:ff:fe:60:37:c8 PTP Port number: 1 PTP interface state: MASTER PTP operational state: UP Forced Master: no Delay request interval(log mean): 0 Announce receipt time out: 3 Announce interval(log mean): -2 Sync interval(log mean): -5 Delay Mechanism: End to End Transport protocol: UDP IPv6 IPv6 Multicast scope ID: Global (0xE) Interface name: Eth1/11 (Po 3) Channel group ID: 3 VRF: default IP Address: 1111:0:0:0:0:0:0:0/64 Port Clock identity: ec:0D:9a:ff:fe:60:37:c8 PTP Port number: 1 PTP interface state: MASTER PTP operational state: UP Forced Master: no Delay request interval(log mean): 0 Announce receipt time out: 3 Announce interval(log mean): -2 Sync interval(log mean): -5 Delay Mechanism: End to End Transport protocol: UDP IPv6 IPv6 Multicast scope ID: Global (0xE)
Related Commands	
Notes	

5.2.6.34 show ptp interface port-channel counters

	show ptp interface port-channel <po-id> counters Displays port statistics on LAG member interfaces.	
Syntax Description	po-id	LAG ID
Default	N/A	
Configuration Mode	Any command mode	
History	3.7.1000	

Example	switch (config) # show ptp interface port-channel 3 counters Eth1/10 RX 0 Sync message count 0 Delay request message count 0 PDelay request message count 0 PDelay response message count 0 Follow Up message count 0 Delay response message count 0 PDelay response follow Up message count 0 Announce message count 0 Signalling message count 0 Management message count TX 0 Sync message count 0 Delay request message count 0 PDelay request message count 0 PDelay response message count 0 Follow Up message count 0 Delay response message count 0 PDelay response follow Up message count 0 Announce message count 0 Signalling message count 1 Management message count 0 Forwarded Management message count Eth1/11 (Po 3) RX 0 Sync message count 0 Delay request message count 0 PDelay request message count 0 PDelay response message count 0 Follow Up message count 0 Delay response message count 0 PDelay response follow Up message count 0 Announce message count 0 Signalling message count 0 Management message count TX 0 Sync message count 0 Delay request message count 0 PDelay request message count 0 PDelay response message count 0 Follow Up message count 0 Delay response message count 0 PDelay response follow Up message count 0 Announce message count 0 Signalling message count 2 Management message count 0 Forwarded Management message count
Related Commands	
Notes	

5.2.6.35 show ptp amt log

	show ptp amt log Displays received GMC clock IDs outside of acceptable master table.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.8100
Example	

```
switch (config) # show ptp amt log

-------------------------------------------------------------------------------
Clock Identity            Interface   VLAN   IP Address     Last Occurrence
-------------------------------------------------------------------------------
04:1D:2D:FF:FE:A5:F3:94   Eth1/2      N/A    192.168.66.7   2018/07/17 19:44:09
03:1D:2D:FF:FE:A5:F3:94   Eth1/2      N/A    192.168.66.7   2018/07/17 19:44:09
```

Related Commands	show ptp amt clear ptp amt log
Notes	

5.2.6.36 show ptp clock

	show ptp clock Displays configuration and operation data of PTP clock.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.4110
Example	switch (config) # show ptp clock Domain: 127 Number of PTP ports: 1 Priority1: 128 Priority2: 128 Clock identity: e41d2d.fffe.46f801 Offset From Master (ns): 65535 Mean path delay (ns): 13303808 Clock Quality Class: 248 Accuracy: 254 Offset (log variance): 65535 Steps Removed from GMC: 1 Local clock time: 13:59:27 Etc/UTC 2017/05/23 ...
Related Commands	
Notes	

5.2.6.37 show ptp clock parent

	show ptp clock parent Displays configuration and operation data of parent PTP clock.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4110	
	3.8.2100	Updated example

Example	switch (config) # show ptp clock parent Parent Clock Parent Clock identity: ec:46:70:ff:fe:0c:e4:82 Parent Port number: 1 GMC GMC Identity: ec:46:70:ff:fe:0c:e4:82 GMC Clock Quality Priority1: 128 Priority2: 128 Class: 6 Accuracy: 33 Offset (log variance): 13563 Time Traceable : 1 (True) Frequency Traceable: 1 (True) PTP Timescale : 1 (True) Time Source : 0x20 (GPS)
Related Commands	
Notes	

5.2.6.38 show ptp forced-master

	show ptp forced-master Displays forced master PTP interfaces.
Syntax Description	N/A
Default	N/A
onfiguration ode	Any command mode
History	3.6.8100
Example	switch (config) # show ptp forced-master --- Port Po VLAN VRF --- Eth1/10 3 N/A default Eth1/11 3 N/A default
Related Commands	show ptp
Notes	

5.2.6.39 show ptp

	show ptp <slot>/<port>[/<subport>] Displays PTP configuration and operation data per Ethernet port.	
Syntax Description	<slot>/<port>/ <subport>	Ethernet port ID (e.g. 1/3/1)
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4110	
	3.6.8100	Updated example
	3.8.2000	Updated example

Example	switch (config) # show ptp 1/1 Interface name: Eth1/1 Channel group ID: N/A VRF: default IP Address: 1111:0:0:0:0:0:0:0/64 Port Clock identity: ec:0D:9a:ff:fe:60:37:c8 PTP Port number: 1 PTP interface state: MASTER Forced Master: no Delay request interval(log mean): 0 Announce receipt time out: 3 Announce interval(log mean): -2 Sync interval(log mean): -3 Delay Mechanism: End to End Transport protocol: UDP IPv6 IPv6 Multicast Scope ID: Global (0xE)
Related Commands	
Notes	

5.2.6.40 show ptp clock foreign-masters

	show ptp clock foreign-masters Displays all PTP foreign masters per each PTP port.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.8.2100
Example	show ptp clock foreign-masters --- Interface Clock-ID P1 P2 CC CA OSLV SR GM --- Eth1/15 ec:46:70:ff:fe:0c:e4:82 128 128 6 33 13563 0 Y Eth1/13 00:80:ea:ff:fe:d0:25:aa 128 1 6 33 20061 0 N
Related Commands	show ptp show log
Notes	

5.2.6.41 show ptp interface ethernet counters

	show ptp interface ethernet <slot>/<port>[/<subport>] counters Displays PTP counters per Ethernet port.	
Syntax Description	<slot>/<port>/ <subport>	Ethernet port ID (e.g. 1/3/1)
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4110	
	3.6.8008	Added VLAN parameter

Example	switch (config) # show ptp interface ethernet 1/5 counters Eth1/5 RX 108 Sync message count 0 Delay request message count 0 PDelay request message count 0 PDelay response message count 108 Follow Up message count 17 Delay response message count 0 PDelay response follow Up message count 54 Announce message count 0 Signaling message count 0 Management message count TX 74188 Sync message count 17 Delay request message count 0 PDelay request message count 0 PDelay response message count 74188 Follow Up message count 0 Delay response message count 0 PDelay response follow Up message count 37117 Announce message count 0 Signaling message count 57 Management message count ...
Related Commands	
Notes	

5.2.6.42 show ptp interface

	show ptp interface Displays PTP configuration and operation data for all PTP-enabled interfaces.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.8.2000

Example	switch (config) # show ptp interface Interface name: Eth1/4 Channel group ID: N/A VRF: default IP Address: 4.4.4.4/24 Port Clock identity: 7c:fe:90:ff:fe:fa:22:08 PTP Port number: 1 PTP interface state: MASTER PTP operational state: UP Forced Master: no Delay request interval(log mean): 0 Announce receipt time out: 3 Announce interval(log mean): -2 Sync interval(log mean): -3 Delay Mechanism: End to End Transport protocol: UDP IPv4 IPv6 Multicast scope ID: N/A Interface name: Eth1/12 (VLAN 12) Channel group ID: 12 VRF: default IP Address: 12.8.8.8/24 Port Clock identity: 7c:fe:90:ff:fe:fa:22:08 PTP Port number: 2 PTP interface state: SLAVE PTP operational state: UP Forced Master: no Delay request interval(log mean): 0 Announce receipt time out: 3 Announce interval(log mean): -2 Sync interval(log mean): -3 Delay Mechanism: End to End Transport protocol: UDP IPv4 IPv6 Multicast scope ID: N/A
Related Commands	show ptp interface ethernet show ptp interface vlan
Notes	

5.2.6.43 show ptp interface ethernet

	show ptp interface ethernet <id> Displays PTP configuration and operation data for the ethernet interface.	
Syntax Description	id	Ethernet ID
Default	N/A	
Configuration Mode	Any command mode	
History	3.8.2000	
Example	switch (config) # show ptp interface ethernet 1/12 Interface name: Eth1/12 (VLAN 12) Channel group ID: 12 VRF: default IP Address: 12.8.8.8/24 Port Clock identity: 7c:fe:90:ff:fe:fa:22:08 PTP Port number: 2 PTP interface state: SLAVE PTP operational state: UP Forced Master: no Delay request interval(log mean): 0 Announce receipt time out: 3 Announce interval(log mean): -2 Sync interval(log mean): -3 Delay Mechanism: End to End Transport protocol: UDP IPv4 IPv6 Multicast scope ID: N/A	

Related Commands	
Notes	

5.2.6.44 show ptp interface vlan

	show ptp interface vlan <vid> Displays PTP configuration and operation data per VLAN.	
Syntax Description	vid	VLAN ID
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	
	3.6.8100	Updated example
	3.8.2000	Updated example

Example	`switch (config) # show ptp interface vlan 1`
	`Interface name: Eth1/15/1 (VLAN 1)`
	`Port Clock identity: 7cfe90.fffe.fa2388`
	`PTP Port number: 1`
	`PTP interface state: SLAVE`
	`PTP operational state: UP`
	`Forced Master: no`
	`Delay request interval(log mean): 0`
	`Announce receipt time out: 3`
	`Announce interval(log mean): -2`
	`Sync interval(log mean): -3`
	`Delay Mechanism: End to End`
	`Transport protocol: UDP IPv6`
	`IPv6 Multicast scope ID: Global (0xE)`	

Related Commands		
Notes		

5.2.6.45 show ptp interface vlan ethernet

	show ptp interface vlan <vid> ethernet <slot>/<port>[/<subport>] Displays PTP configuration and operation data for specified VLAN member interface for a specified Ethernet port.	
Syntax Description	vid	VLAN ID
	<slot>/<port>/<subport>	Ethernet port ID (e.g. 1/3/1)
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	
	3.8.2000	Updated example

Example	switch (config) # show ptp interface vlan 1 ethernet 1/15/1 Interface name: Eth1/15/1 (VLAN 1) Port Clock identity: 7cfe90.fffe.fa2388 PTP Port number: 1 PTP interface state: FAULTY PTP operational state: UP Delay request interval(log mean): 0 Announce receipt time out: 3 Announce interval(log mean): -2 Sync interval(log mean): -3 Delay Mechanism: End to End Transport protocol: UDP IPv4 IPv6 Multicast scope ID: N/A
Related Commands	
Notes	

5.2.6.46 show ptp interface vlan counters

	show ptp interface vlan <vid> counters Displays PTP counters per VLAN.	
Syntax Description	vid	VLAN ID
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	
	3.8.2000	Added example
Example	switch (config) # show ptp interface vlan 3 counters Eth1/3 (VLAN 3) RX 0 Sync message count 0 Delay request message count 0 PDelay request message count 0 PDelay response message count 0 Follow Up message count 0 Delay response message count 0 PDelay response follow Up message count 0 Announce message count 0 Signalling message count 0 Management message count TX 19851 Sync message count 0 Delay request message count 0 PDelay request message count 0 PDelay response message count 19851 Follow Up message count 0 Delay response message count 0 PDelay response follow Up message count 9928 Announce message count 0 Signalling message count 2 Management message count 0 Forwarded Management message count	
Related Commands		
Notes		

5.2.6.47 show ptp interface vlan ethernet counters

	show ptp interface vlan <vid> ethernet <slot>/<port>[/<subport>] counters Displays PTP counters per VLAN for a specified Ethernet port.

Syntax Description	vid	VLAN ID
	<slot>/<port>/<subport>	Ethernet port ID (e.g. 1/3/1)
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	
Example	switch (config) # show ptp interface vlan 1 ethernet 1/15/1 counters Eth1/15/1 (VLAN 1) RX 0 Sync message count 0 Delay request message count 0 PDelay request message count 0 PDelay response message count 0 Follow Up message count 0 Delay response message count 0 PDelay response follow Up message count 0 Announce message count 0 Signaling message count 0 Management message count TX 0 Sync message count 0 Delay request message count 0 PDelay request message count 0 PDelay response message count 0 Follow Up message count 0 Delay response message count 0 PDelay response follow Up message count 0 Announce message count 0 Signaling message count 0 Management message count	
Related Commands		
Notes		

5.2.6.48 show ptp time-property

	show ptp time-property Displays PTP time-property parameters (time source, current utc offset etc).
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.8.2100
Example	switch (config) # show ptp time-property Current UTC Offset valid: 1 (True) Current UTC Offset : 37 Leap59 : 0 (False) Leap61 : 0 (False) Time Traceable : 1 (True) Frequency Traceable : 1 (True) PTP Timescale : 1 (True) Time Source : 0x20 (GPS)
Related Commands	
Notes	

5.2.6.49 show ptp status

	show ptp status Displays the last 100 entries for Offset from Master and Mean Path Delay values.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any configuration mode
History	3.8.2100
Example	switch (config) # show ptp status PTP mode : Boundary Clock PTP Offset Threshold (ns) : -100000, 100000 PTP Mean Path Delay Threshold (ns): 1000000000 --- Interface Time Offset from Master (ns) Mean Path Delay (ns) --- Eth1/15 2019/11/13 16:32:00.774 -21 424 Eth1/15 2019/11/13 16:32:00.649 -28 424 Eth1/15 2019/11/13 16:32:00.524 18 424 Eth1/15 2019/11/13 16:32:00.399 6 424 Eth1/15 2019/11/13 16:32:00.274 28 423 Eth1/15 2019/11/13 16:32:00.149 -16 424 Eth1/15 2019/11/13 16:32:00.025 -7 425 Eth1/15 2019/11/13 16:31:59.899 17 425 Eth1/15 2019/11/13 16:31:59.775 9 422 Eth1/15 2019/11/13 16:31:59.650 -3 420 Eth1/15 2019/11/13 16:31:59.525 -16 425 Eth1/15 2019/11/13 16:31:59.400 -23 422 Eth1/15 2019/11/13 16:31:59.275 17 422
Related Commands	
Notes	

5.2.6.50 PTP Debuggability Logging Examples

5.2.6.50.1 Change of the State of Particular PTP Port

```
Nov 11 15:33:09 arc-switch142 ptp4l: [351.341] PTP [Debuggability]: PTP Grandmaster clock has changed
from 000000.0000.000000 to ec0d9a.fffe.603848
Nov 11 15:33:09 arc-switch142 ptp4l: [351.341] port 0: hybrid_e2e only works with E2E
Nov 11 15:33:09 arc-switch142 ptp4l: [351.342] port 1: Interface Eth1/10 state changed from INITIALIZING to
LISTENING on INIT_COMPLETE
Nov 11 15:33:09 arc-switch142 ptp4l: [351.342] port 0: Interface  state changed from INITIALIZING to
LISTENING on INIT_COMPLETE
Nov 11 15:33:09 arc-switch142 ptp4l: [351.342] port 1: link down
Nov 11 15:33:09 arc-switch142 ptp4l: [351.342] port 1: Interface Eth1/10 state changed from LISTENING to
FAULTY on FAULT_DETECTED (FT_UNSPECIFIED)
Nov 11 15:33:09 arc-switch142 ptp4l: [351.343] selected local clock ec0d9a.fffe.603848 as best master
Nov 11 15:33:09 arc-switch142 ptp4l: [351.343] assuming the grand master role
Nov 11 15:33:09 arc-switch142 ptp4l: [351.343] PTP [Debuggability]: PTP Grandmaster clock has changed
from ec0d9a.fffe.603848 to ec0d9a.fffe.603848
Nov 11 15:33:09 arc-switch142 pm[4868]: [pm.NOTICE]: Launched phc2sys (PTP phc2sys daemon) with pid
7870
Nov 11 15:33:09 arc-switch142 ptp4l: [351.455] port 1: link up
Nov 11 15:33:09 arc-switch142 ptp4l: [351.456] port 1: Interface Eth1/10 state changed from FAULTY to
LISTENING on INIT_COMPLETE
Nov 11 15:33:10 arc-switch142 ptp4l: [352.295] PTP [Debuggability]: Matched Announce interval on Eth1/10.
Configured -2, Received -2
Nov 11 15:33:10 arc-switch142 ptp4l: [352.295] port 1: new foreign master ec0d9a.fffe.6037c8-1
Nov 11 15:33:10 arc-switch142 ptp4l: [352.402] port 1: Interface Eth1/10 state changed from LISTENING to
MASTER on ANNOUNCE_RECEIPT_TIMEOUT_EXPIRES
Nov 11 15:33:10 arc-switch142 ptp4l: [352.402] selected local clock ec0d9a.fffe.603848 as best master
Nov 11 15:33:10 arc-switch142 ptp4l: [352.402] assuming the grand master role
Nov 11 15:33:10 arc-switch142 ptp4l: [352.402] PTP [Debuggability]: PTP Grandmaster clock has changed
from ec0d9a.fffe.603848 to ec0d9a.fffe.603848
Nov 11 15:33:10 arc-switch142 ptp4l: [352.419] PTP [Debuggability]: Matched Sync interval on Eth1/10.
Configured -3, Received -3
Nov 11 15:33:11 arc-switch142 ptp4l: [352.795] selected best master clock ec0d9a.fffe.6037c8
Nov 11 15:33:11 arc-switch142 ptp4l: [352.795] PTP [Debuggability]: PTP Grandmaster clock has changed
from ec0d9a.fffe.603848 to ec0d9a.fffe.6037c8
Nov 11 15:33:11 arc-switch142 ptp4l: [352.795] port 1: Interface Eth1/10 state changed from MASTER to
UNCALIBRATED on RS_SLAVE
Nov 11 15:33:11 arc-switch142 ptp4l: [353.044] PTP slave port Eth1/10 High offset from Master 635155 (ns)
Nov 11 15:33:11 arc-switch142 ptp4l: [353.169] PTP slave port Eth1/10 High offset from Master 635353 (ns)
Nov 11 15:33:11 arc-switch142 ptp4l: [353.294] port 1: Interface Eth1/10 state changed from UNCALIBRATED
to SLAVE on MASTER_CLOCK_SELECTED
```

5.2.6.50.2 Change of Grandmaster Clock

```
Nov 11 15:33:09 arc-switch142 ptp4l: [351.341] PTP [Debuggability]: PTP Grandmaster clock has changed
from 000000.0000.000000 to ec0d9a.fffe.603848
Nov 11 15:33:09 arc-switch142 ptp4l: [351.341] port 0: hybrid_e2e only works with E2E
Nov 11 15:33:09 arc-switch142 ptp4l: [351.342] port 1: Interface Eth1/10 state changed from INITIALIZING to
LISTENING on INIT_COMPLETE
Nov 11 15:33:09 arc-switch142 ptp4l: [351.342] port 0: Interface  state changed from INITIALIZING to
LISTENING on INIT_COMPLETE
Nov 11 15:33:09 arc-switch142 ptp4l: [351.342] port 1: link down
Nov 11 15:33:09 arc-switch142 ptp4l: [351.342] port 1: Interface Eth1/10 state changed from LISTENING to
FAULTY on FAULT_DETECTED (FT_UNSPECIFIED)
Nov 11 15:33:09 arc-switch142 ptp4l: [351.343] selected local clock ec0d9a.fffe.603848 as best master
Nov 11 15:33:09 arc-switch142 ptp4l: [351.343] assuming the grand master role
Nov 11 15:33:09 arc-switch142 ptp4l: [351.343] PTP [Debuggability]: PTP Grandmaster clock has changed
from ec0d9a.fffe.603848 to ec0d9a.fffe.603848
Nov 11 15:33:09 arc-switch142 pm[4868]: [pm.NOTICE]: Launched phc2sys (PTP phc2sys daemon) with pid
7870
Nov 11 15:33:09 arc-switch142 ptp4l: [351.455] port 1: link up
Nov 11 15:33:09 arc-switch142 ptp4l: [351.456] port 1: Interface Eth1/10 state changed from FAULTY to
LISTENING on INIT_COMPLETE
Nov 11 15:33:10 arc-switch142 ptp4l: [352.295] PTP [Debuggability]: Matched Announce interval on Eth1/10.
Configured -2, Received -2
Nov 11 15:33:10 arc-switch142 ptp4l: [352.295] port 1: new foreign master ec0d9a.fffe.6037c8-1
Nov 11 15:33:10 arc-switch142 ptp4l: [352.402] port 1: Interface Eth1/10 state changed from LISTENING to
MASTER on ANNOUNCE_RECEIPT_TIMEOUT_EXPIRES
Nov 11 15:33:10 arc-switch142 ptp4l: [352.402] selected local clock ec0d9a.fffe.603848 as best master
Nov 11 15:33:10 arc-switch142 ptp4l: [352.402] assuming the grand master role
Nov 11 15:33:10 arc-switch142 ptp4l: [352.402] PTP [Debuggability]: PTP Grandmaster clock has changed
from ec0d9a.fffe.603848 to ec0d9a.fffe.603848
Nov 11 15:33:10 arc-switch142 ptp4l: [352.419] PTP [Debuggability]: Matched Sync interval on Eth1/10.
Configured -3, Received -3
Nov 11 15:33:11 arc-switch142 ptp4l: [352.795] selected best master clock ec0d9a.fffe.6037c8
Nov 11 15:33:11 arc-switch142 ptp4l: [352.795] PTP [Debuggability]: PTP Grandmaster clock has changed
from ec0d9a.fffe.603848 to ec0d9a.fffe.6037c8
Nov 11 15:33:11 arc-switch142 ptp4l: [352.795] port 1: Interface Eth1/10 state changed from MASTER to
UNCALIBRATED on RS_SLAVE
Nov 11 15:33:11 arc-switch142 ptp4l: [353.044] PTP slave port Eth1/10 High offset from Master 635155 (ns)
Nov 11 15:33:11 arc-switch142 ptp4l: [353.169] PTP slave port Eth1/10 High offset from Master 635353 (ns)
Nov 11 15:33:11 arc-switch142 ptp4l: [353.294] port 1: Interface Eth1/10 state changed from UNCALIBRATED
to SLAVE on MASTER_CLOCK_SELECTED
```

Announce Interval Mismatch Notification

```
Nov 11 15:41:10 arc-switch142 ptp4l: [869.220] PTP [Debuggability]: PTP Grandmaster clock has changed
from 000000.0000.000000 to ec0d9a.fffe.603848
Nov 11 15:41:10 arc-switch142 ptp4l: [869.221] port 0: hybrid_e2e only works with E2E
Nov 11 15:41:10 arc-switch142 ptp4l: [869.221] port 1: Interface Eth1/10 state changed from INITIALIZING to
LISTENING on INIT_COMPLETE
Nov 11 15:41:10 arc-switch142 ptp4l: [869.221] port 0: Interface  state changed from INITIALIZING to
LISTENING on INIT_COMPLETE
Nov 11 15:41:10 arc-switch142 pm[4868]: [pm.NOTICE]: Launched phc2sys (PTP phc2sys daemon) with pid
8918
Nov 11 15:41:10 arc-switch142 ptp4l: [869.284] PTP [Debuggability]: Matched Sync interval on Eth1/10.
Configured -3, Received -3
Nov 11 15:41:10 arc-switch142 ptp4l: [869.284] PTP [Debuggability]: Mismatch Announce interval on Eth1/10.
Configured -1, Received -3
Nov 11 15:41:10 arc-switch142 ptp4l: [869.284] port 1: new foreign master ec0d9a.fffe.6037c8-1
Nov 11 15:41:10 arc-switch142 ptp4l: [869.534] selected best master clock ec0d9a.fffe.6037c8
Nov 11 15:41:10 arc-switch142 ptp4l: [869.534] PTP [Debuggability]: PTP Grandmaster clock has changed
from ec0d9a.fffe.603848 to ec0d9a.fffe.6037c8
Nov 11 15:41:10 arc-switch142 ptp4l: [869.534] port 1: Interface Eth1/10 state changed from LISTENING to
UNCALIBRATED on RS_SLAVE
Nov 11 15:41:11 arc-switch142 ptp4l: [869.909] port 1: Interface Eth1/10 state changed from UNCALIBRATED
to SLAVE on MASTER_CLOCK_SELECTED
Nov 11 15:42:34 arc-switch142 ptp4l: [953.018] PTP [Debuggability]: Matched Announce interval on Eth1/10.
Configured -1, Received -1
```

Sync Interval Mismatch Notification

```
Nov 11 16:05:34 arc-switch142 ptp4l: [2332.929] PTP [Debuggability]: PTP Grandmaster clock has changed
from ec0d9a.fffe.603848 to ec0d9a.fffe.6037c8
Nov 11 16:05:34 arc-switch142 ptp4l: [2332.929] port 1: Interface Eth1/10 state changed from MASTER to
UNCALIBRATED on RS_SLAVE
Nov 11 16:05:34 arc-switch142 ptp4l: [2333.053] PTP [Debuggability]: Mismatch Sync interval on Eth1/10.
Configured -3, Received -2
Nov 11 16:05:34 arc-switch142 ptp4l: [2333.303] port 1: Interface Eth1/10 state changed from UNCALIBRATED
to SLAVE on MASTER_CLOCK_SELECTED
Nov 11 16:06:14 arc-switch142 ptp4l: [2372.799] port 1: Interface Eth1/10 state changed from SLAVE to
MASTER on ANNOUNCE_RECEIPT_TIMEOUT_EXPIRES
Nov 11 16:06:14 arc-switch142 ptp4l: [2372.799] selected local clock ec0d9a.fffe.603848 as best master
Nov 11 16:06:14 arc-switch142 ptp4l: [2372.799] assuming the grand master role
Nov 11 16:06:14 arc-switch142 ptp4l: [2372.799] PTP [Debuggability]: PTP Grandmaster clock has changed
from ec0d9a.fffe.6037c8 to ec0d9a.fffe.603848
Nov 11 16:06:14 arc-switch142 ptp4l: [2372.943] selected best master clock ec0d9a.fffe.6037c8
Nov 11 16:06:14 arc-switch142 ptp4l: [2372.943] PTP [Debuggability]: PTP Grandmaster clock has changed
from ec0d9a.fffe.603848 to ec0d9a.fffe.6037c8
Nov 11 16:06:14 arc-switch142 ptp4l: [2372.943] port 1: Interface Eth1/10 state changed from MASTER to
UNCALIBRATED on RS_SLAVE
Nov 11 16:06:14 arc-switch142 ptp4l: [2373.317] PTP [Debuggability]: Mismatch Sync interval on Eth1/10.
Configured -3, Received -1
Nov 11 16:06:15 arc-switch142 ptp4l: [2373.817] port 1: Interface Eth1/10 state changed from UNCALIBRATED
to SLAVE on MASTER_CLOCK_SELECTED
Nov 11 16:06:33 arc-switch142 ptp4l: [2392.739] port 1: Interface Eth1/10 state changed from SLAVE to
MASTER on ANNOUNCE_RECEIPT_TIMEOUT_EXPIRES
Nov 11 16:06:33 arc-switch142 ptp4l: [2392.739] selected local clock ec0d9a.fffe.603848 as best master
Nov 11 16:06:33 arc-switch142 ptp4l: [2392.739] assuming the grand master role
Nov 11 16:06:33 arc-switch142 ptp4l: [2392.739] PTP [Debuggability]: PTP Grandmaster clock has changed
from ec0d9a.fffe.6037c8 to ec0d9a.fffe.603848
Nov 11 16:06:34 arc-switch142 ptp4l: [2392.978] PTP [Debuggability]: Matched Sync interval on Eth1/10.
Configured -3, Received -3
Nov 11 16:06:34 arc-switch142 ptp4l: [2392.979] selected best master clock ec0d9a.fffe.6037c8
Nov 11 16:06:34 arc-switch142 ptp4l: [2392.979] PTP [Debuggability]: PTP Grandmaster clock has changed
from ec0d9a.fffe.603848 to ec0d9a.fffe.6037c8
```

5.3 Replace CRC with Timestamp

> Replacing CRC field with a timestamp is only supported on Spectrum-2 and Spectrum-3 systems.

In some applications, it is important to know the exact time when a packet arrived at the switch in order to analyze networkwide application behavior. In order to achieve this capability, it is possible to mark the packet that leaves the switch with a timestamp that shows when this packet arrived.

One of the use cases is a mirroring setup when an original packet is forwarded by the system, but its mirrored copy is sent to a collector for analysis together with a timestamp that will help analyzer to rebuild the sequence of events in the network.

5.3.1 Main Functionality

The feature gives a possibility to configure the following functionality:

- Disabling checking of the CRC on the ingress port
- Disabling recalculation of the CRC on the egress port
- Disabling replacement of the FCS (Frame Check Sequence) field in the packet with a timestamp globally

CRC checking operation is enabled by default and is performed for incoming packets. Disabling CRC checking is required in cases when we want the packet which has timestamp instead of FCS field to traverse the switch through the ingress port. Otherwise (when CRC checking is enabled), it will be dropped as the packet that has bad CRC value.

CRC recalculation operation is done on the outgoing packets by default. Disabling CRC recalculation is required in cases when we want the packet to pass through the egress port and preserve the timestamp in the FCS field. Otherwise (when CRC recalculation is enabled), the packet's FCS field will be overwritten by actual CRC value during the recalculation process.

A timestamp that is placed into FCS field in the packet can be obtained from the following time sources according to priority:

- NTP (if running)
- Local clock

The timestamp identifies a time when the packet is entered into the system. It is presented in UTC format and overwrites 26 bits of the FCS field as follows:

- 24 bits in nanoseconds [29:6]
- 2 bits in seconds [31:30]

Replacing the FCS field with a timestamp is enabled globally in NVIDIA Onyx by default. Despite this, packets will still leave the switch without the timestamp – it will be overwritten on the egress port during recalculation process (unless the CRC recalculation is disabled by the user).

5.3.2 Setup Configuration

In order to ensure that the timestamp will traverse through the switch, the following configuration should be applied:

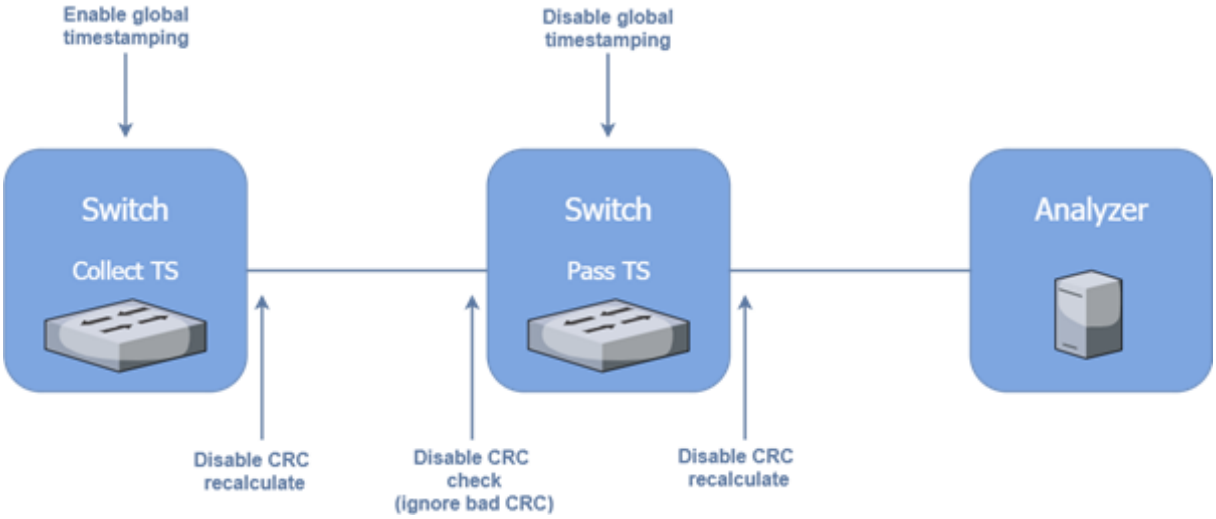

The first switch is collecting timestamp – the timestamp will be set when a packet entered the system through the ingress port. In order to preserve the timestamp, the CRC recalculation should be disabled on the egress port.

The packet with a timestamp should pass through the second switch transparently. For this purpose, both CRC check and recalculation must be disabled on the ingress and egress ports accordingly.

To disable CRC recalculation on the port:

1. Log in as admin.
2. Enter config mode. Run:

```
switch > enable
switch # configure terminal
```

3. Disable recalculation of the CRC on the specific egress port. Run:

```
switch (config interface ethernet 1/1) # fcs egress disable-recalculate
```

To disable CRC check on the port:

1. Log in as admin.
2. Enter config mode. Run:

```
switch > enable
switch # configure terminal
```

3. Disable checking of the CRC on the specific ingress port. Run:

```
switch (config interface ethernet 1/1) # fcs ingress disable-check
```

To disable timestamping globally:

1. Log in as admin.
2. Enter config mode. Run:

```
switch > enable
switch # configure terminal
```

3. Disable timestamping in the system. Run:

```
switch (config) # system timestamp disable
```

5.3.3 Replace CRC with Timestamp Commands

- 5.3.1 Main Functionality
- 5.3.2 Setup Configuration
- 5.3.3 Replace CRC with Timestamp Commands
  - 5.3.3.1 fcs igress disable-check
  - 5.3.3.2 fcs egress disable-recalculate
  - 5.3.3.3 system timestamp disable

5.3.3.1 fcs igress disable-check

	fcs ingress disable-check no fcs ingress disable-check Disables checking of the CRC value in the ingress packets received on the interface. The “no” form of the command enables checking of the CRC value.
Syntax Description	N/A
Default	Enabled CRC
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.9.1000
Example	`switch (config) # interface ethernet 1/1 fcs ingress disable-check`
Related Commands	fcs egress disable-recalculate show interfaces ethernet show interfaces port-channel show interfaces mlag-port-channel system timestamp disable
Notes	Disable CRC check cannot be configured on the LAG or MLAG ethernet members alone.

5.3.3.2 fcs egress disable-recalculate

	fcs egress disable-recalculate no fcs egress disable-recalculate Disables recalculation of the CRC value in the egress packets being sent from the interface. The “no” form of the command enables recalculation of the CRC value.
Syntax Description	N/A
Default	Enabled CRC recalculation

Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.9.1000
Example	`switch (config) # interface ethernet 1/1 fcs egress disable-recalculate`
Related Commands	fcs ingress disable-check show interfaces ethernet show interfaces port-channel show interfaces mlag-port-channel system timestamp disable
Notes	Disable CRC recalculate cannot be configured on the LAG or MLAG ethernet members alone.

5.3.3.3 system timestamp disable

	system timestamp disable no system timestamp disable Disables replacement of the CRC/FCS field in the packet with a timestamp in the system. The "no" form of the command enables replacement of the CRC/FCS field with a timestamp.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config
History	3.9.1000
Example	`switch (config) # system timestamp disable`
Related Commands	fcs ingress disable-check fcs egress disable-recalculate show interfaces ethernet show interfaces port-channel show interfaces mlag-port-channel
Notes	Timestamping is enabled in the system by default. Despite this, packets will still leave the switch without the timestamp—it will be overwritten on the egress port during recalculation process.

6 Network Management Interfaces

6.1 SNMP

Simple Network Management Protocol (SNMP), is a network protocol for the management of a network and the monitoring of network devices and their functions. SNMP supports asynchronous event (trap) notifications and queries.NVIDIA Onyx supports:

- SNMP versions v1, v2c and v3
- SNMP trap notifications
- Standard MIBs
- Private MIBs

6.1.1 Standard MIBs

The following table presents the supported textual conventions and conformance MIBs:

MIB	Standard
INET-ADDRESS-MIB	RFC-4001
SNMPV2-CONF	
SNMPV2-TC	RFC 2579
SNMPV2-TM	RFC 3417
SNMP-USM-AES-MIB	RFC 3826
IANA-LANGUAGE-MIB	RFC 2591
IANA-RTPROTO-MIB	RFC 2932
IANAifType-MIB	
IANA-ADDRESS-FAMILY-NUMBERS-MIB	
IGMP-STD-MIB	RFC2933 (See IGMP-STD-MIB Information section)

The following table presents the supported chassis and switch MIBs:

6.1.2 Private MIBs

MIB	Description
MELLANOX-SMI-MIB	Private MIB main structure (no objects)
MELLANOX-PRODUCTS-MIB	List of OID - per managed system (sysObjID)
MELLANOX-IF-VPI-MIB	IfTable extensions
MELLANOX-EFM-MIB	Partially deprecated MIB (based on Mellanox-MIB) Traps definitions and test trap set scalar are supported.
MELLANOX-ENTITY-MIB	Enhances the standard ENTITY-MIB (contains GUID and ASIC revision).
MELLANOX-POWER-CYCLE	Allows rebooting the switch system

MIB	Description
MELLANOX-SW-UPDATE-MIB	Allows viewing what SW images are installed, uploading and installing new SW images
MELLANOX-CONFIG-DB	Allows loading, uploading, or deleting configuration files
MELLANOX-ENTITY-STATE-MIB	Extension to support state change traps Note: Currently supported for power supply insertion and extraction only
MELLANOX-XSTP-MIB	Extension to support STP information
MELLANOX-DCB-TRAPS	Extension traps for ETC and PFC
MELLANOX-QOS	Proprietary QoS MIBs
MELLANOX-WJH-MIB	Defines what-just-happened traps

Private MIBs can be downloaded from the support website.

6.1.3 Proprietary Traps

The following private traps are supported by the NVIDIA Onyx

MELLANOX-EFM-MIB:

Trap	Action Required
asicChipDown	Reboot the system.
asicOverTempReset	Check fans and environmental temperature.
asicOverTemp	Check fans and environmental temperature.
lowPower	Add/connect power supplies.
internalBusError	N/A
procCrash	Generate SysDump and contact support.
cpuUtilHigh	N/A
procUnexpectedExit	Generate SysDump and contact support.
diskSpaceLow	Clean images and sysDump files using the commands “image delete” and “file debug-dump delete”.
systemHealthStatus	Refer to Health Status table.
lowPowerRecover	N/A
insufficientFans	Check Fans and environmental conditions.
insufficientFansRecover	N/A
insufficientPower	Add/connect power supplies, or change power mode using the command “power redundancy mode”.
insufficientPowerRecover	N/A

For additional information refer to MELLANOX-EFM-MIB.

For event-to-MIB mapping, please refer to "Supported Event Notifications and MIB Mapping".

The only MELLANOX-POWER-CYCLE trap supported is mellanoxPowerCyclePlannedReload.

6.1.4 Configuring SNMP

Activate the SNMP server on your switch by running:

```
switch (config) # snmp-server enable
switch (config) # snmp-server enable notify
switch (config) # snmp-server community public ro
switch (config) # snmp-server contact "contact name"
switch (config) # snmp-server host <host IP address> traps version 2c public
switch (config) # snmp-server location "location name"
switch (config) # snmp-server user admin v3 enable
switch (config) # snmp-server user admin v3 prompt auth md5 priv des
```

Community strings are case sensitive.

6.1.5 Resetting SNMPv3 Engine ID

Switch systems shipped with an OS versions older than 3.6.6102 have all had the exact same SNMPv3 engine ID. Going forward, however, all switch systems will ship with a system-specific engine ID.

Upgrading the OS version to 3.6.6102 or higher does not automatically change the current engine ID. That can be done through one of the following methods after performing the software upgrade:

- Changing a switch system's profile
- Running "reset factory"
- Using the command "snmp-server engineID reset" (for more details, please see the procedure below)

To reset SNMP engine ID using "snmp-server engineID reset":

Prerequisites:

If any of the following SNMP configurations exist, please delete/disable them and re-enable/ reconfigure them only after SNMP engine ID reset is performed:

1. Make sure SNMP is disabled. Run:

   ```
   switch (config) # no snmp-server enable
   ```

2. Make sure no SNMP trap host is configured. Run:

   ```
   switch (config) # no snmp-server host <ip-address>
   ```

3. Make sure no SNMP users are configured. Run:

```
switch (config) # no snmp-server user <username> v3
```

Procedure:

1. Check existing engine ID:

```
switch (config) # show snmp engineID
Local SNMP engineID: <current_key>
```

2. Reset existing engine ID:

```
switch (config) # snmp-server engineID reset
```

3. Verify new engine ID:

```
switch (config) # show snmp engineID
Local SNMP engineID: <new_key>
```

6.1.6 Configuring an SNMPv3 User

To configure an SNMPv3 user:

1. Configure the user using the command:

```
switch (config) # snmp-server user [role] v3 prompt auth <hash type> priv <privacy type>
```

   Where:
   - user role—admin
   - auth type—md5 or sha or sha224 or sha256 or sha384 or sha512
   - priv type—des or aes-128 or 3des or aes-192 or aes-256 or aes-192-cfb or aes-256-cfb
2. Enter authentication password and its confirmation.
3. Enter privacy password and its confirmation:

```
switch (config) # snmp-server user admin v3 prompt auth md5 priv des
Auth password: ********
      Confirm: ********
Privacy password: ********
         Confirm: ********
```

   To retrieve the system table, run the following SNMP command:

```
snmpwalk -v3 -l authPriv -a MD5 -u admin -A "<Authentication password>" -x DES -X "<privacy password>"
<system ip> SNMPv2-MIB::system
```

6.1.7 Configuring SNMP Notifications (Traps or Informs)

1. Make sure SNMP and SNMP notification are enable. Run:

```
switch (config) # snmp-server enable
switch (config) # snmp-server enable notify
```

2. Configure SNMP host with the desired arguments (IP Address, SNMP version, authentication methods). More than one host can be configured. Each host may have different attributes. Run:

```
switch (config) # snmp-server host 10.134.47.3 traps version 3 user my-username auth sha my-password
```

3. Verify the SNMP host configuration. Run:

```
switch (config) # show snmp host
Notifications enabled:            yes
Default notification community:   public
Default notification port:        162

Notification sinks:

   10.134.47.3
      Enabled:                  yes
      Port:                     162 (default)
      Notification type:        SNMP v3 trap
      Username:                 my-username
      Authentication type:      sha
      Privacy type:             aes-128
      Authentication password: (set)
      Privacy password:         (set)
```

4. Configure the desired event to be sent via SNMP. Run:

```
switch (config) # snmp-server notify event interface-up
```

> This particular event is used as an example only.

5. Verify the list of traps and informs being sent to out of the system. Run:

```
switch (config) # show snmp events
Events for which traps will be sent:
  asic-chip-down: ASIC (Chip) Down
  cpu-util-high: CPU utilization has risen too high
  disk-space-low: Filesystem free space has fallen too low
  health-module-status: Health module Status
  insufficient-fans: Insufficient amount of fans in system
  insufficient-fans-recover: Insufficient amount of fans in system recovered
  insufficient-power: Insufficient power supply
  interface-down: An interface's link state has changed to down
  interface-up: An interface's link state has changed to up
  internal-bus-error: Internal bus (I2C) Error
  liveness-failure: A process in the system was detected as hung
  low-power: Low power supply
  low-power-recover: Low power supply Recover
  new_root: local bridge became a root bridge
  paging-high: Paging activity has risen too high
  power-redundancy-mismatch: Power redundancy mismatch
  process-crash: A process in the system has crashed
  process-exit: A process in the system unexpectedly exited
  snmp-authtrap: An SNMP v3 request has failed authentication
  topology_change: local bridge triggered a topology change
  unexpected-shutdown: Unexpected system shutdown
```

> To print event notifications to the terminal (SSH or CONSOLE) refer to “Monitor”.

> For the SNMPv1 traps or informs, by default, the "agent address" field is set to the IP address of the "mgmt0" interface. In the case that "source interface" is configured to the same VRF which is used for SNMPv1 traps or informs, the IP address of the source interface is used for "agent address" field. In other cases (e.g., if source interface might be configured in some other VRF), "127.0.0.1" is used for the "agent address".

6.1.8 SNMP SET Operations

The OS allows the user to use SET operations via SNMP interface. This is needed to configure a user/community supporting SET operations.

6.1.8.1 Enabling SNMP SET

To allow SNMP SET operations using SNMPv1/v2:

1. Enable SNMP communities. Run:

```
switch (config) # snmp-server enable communities
```

2. Configure a read-write community. Run:

```
switch (config) # snmp-server community my-community-name rw
```

3. Make sure SNMP communities are enabled (they are enabled by default). Make sure "(DISABLED)" does not appear beside "Read-only communities" / "Read-write communities". Run:

```
switch (config) # show snmp

SNMP enabled   : yes
SNMP port      : 161
System contact :
System location:

Read-only communities:
  public

Read-write communities:
  my-community-name

Interface listen enabled: yes

Listen Interfaces:
  Interface: mgmt0

switch (config) # show snmp
No Listen Interfaces.
```

4. Configure this RW community in your MIB browser.

To allow SNMP SET operations using SNMPv3:

1. Create an SNMPv3 user. Run:

```
switch (config) # snmp-server user myuser v3 auth sha <password1> priv aes-128 <password2>
```

> It is possible to use other configuration options not specified in the example above. Please refer to the command "snmp-server user" for more information.

2. Make sure the username is enabled for SET access and has admin capability level. Run:

```
switch (config) # show snmp user
User name: myuser
   Enabled overall:          yes
   Authentication type:      sha
   Privacy type:             aes-128
   Authentication password: (set)
   Privacy password:         (set)
   Require privacy:          yes
   SET access:
      Enabled:               yes
      Capability level:      admin
```

The OS supports the OIDs for SET operation listed in the following table which are expanded upon in the following subsections.

	OID Name	OID
MELLANOX-EFM-MIB	sendTestTrapSet	1.3.6.1.4.1.33049.2.1.1.1.6.0
SNMPv2-MIB	sysName	1.3.6.1.2.1.1.5.0
MELLANOX-CONFIG-DB	mellanoxConfigDBCmdExecute mellanoxConfigDBCmdFilename mellanoxConfigDBCmdStatus mellanoxConfigDBCmdStatusString mellanoxConfigDBCmdUri	1.3.6.1.4.1.33049.12.1.1.2.3.0 1.3.6.1.4.1.33049.12.1.1.2.2.0 1.3.6.1.4.1.33049.12.1.1.2.4.0 1.3.6.1.4.1.33049.12.1.1.2.5.0 1.3.6.1.4.1.33049.12.1.1.2.1.0
MELLANOX-POWER-CYCLE	mellanoxPowerCycleCmdExecute mellanoxPowerCycleCmdStatus mellanoxPowerCycleCmdStatusString	1.3.6.1.4.1.33049.10.1.1.2.1.0 1.3.6.1.4.1.33049.10.1.1.2.2.0 1.3.6.1.4.1.33049.10.1.1.2.3.0
MELLANOX-SW-UPDATE	mellanoxSWUpdateCmdSetNext mellanoxSWUpdateCmdUri mellanoxSWUpdateCmdExecute mellanoxSWUpdateCmdStatus mellanoxSWUpdateCmdStatusString mellanoxSWActivePartition mellanoxSWNextBootPartition	1.3.6.1.4.1.33049.11.1.1.2.1.0 1.3.6.1.4.1.33049.11.1.1.2.2.0 1.3.6.1.4.1.33049.11.1.1.2.3.0 1.3.6.1.4.1.33049.11.1.1.2.4.0 1.3.6.1.4.1.33049.11.1.1.2.5.0 1.3.6.1.4.1.33049.11.1.1.3.0.0 1.3.6.1.4.1.33049.11.1.1.4.0.0

6.1.8.2 Sending a Test Trap SET Request

The OS allows the user to use test the notification mechanism via SNMP SET. Sending a SET request with the designated OID triggers a test trap.

Prerequisites:

1. Enable SET operations by following the instructions in “Enabling SNMP SET”.
2. Configure host to which to send SNMP notifications.
3. Set a trap receiver in the MIB browser.

Procedure:

1. Send a SET request to the switch IP with the OID 1.3.6.1.4.1.33049.2.1.1.1.6.0.
2. Make sure the test trap is received by the aforementioned trap receiver (OID: 1.3.6.1.4.1.33049.2.1.2.13).

6.1.8.3 Setting Hostname with SNMP

The OS supports setting system hostname using an SNMP SET request as described in SNMPv2-MIB (sysName, OID: 1.3.6.1.2.1.1.5.0).

The restrictions on setting a hostname via CLI also apply to setting a hostname through SNMP. Refer to the command “hostname” for more information.

6.1.8.4 Power Cycle with SNMP

The OS supports power cycling its systems using an SNMP SET request as described in MELLANOX-POWER-CYCLE MIB.

Power cycle command is issued via the OID mellanoxPowerCycleCmdExecute. The following options are available:

- Reload—saves any unsaved configuration and reloads the switch
- Reload discard—reboots the system and discards of any unsaved changes
- Reload force—forces an expedited reload on the system even if it is busy without saving unsaved configuration (equals the CLI command reload force)

6.1.8.5 Changing Configuration with SNMP

The OS supports making configuration changes on its systems using SNMP SET requests. Configuration requests are performed by setting several values (arguments) and then executing a command by setting the value for the relevant operation.

It is possible to set the parameters and execute the commands on the same SNMP request or separate them to several SET operations. Upon executing a command, the values of its arguments remain and can be read using GET commands.

Once a command is executed there may be two types of errors:

- Immediate: This error results in a failure of the SNMP request. This means a critical error in the SNMP request has occurred or that a previous SET request is being executed
- Delayed: The SET request has been accepted by the switch but an error occurred during its execution.

For example, when performing a fetch (download) operation, an immediate error can occur when the given URL is invalid. A delayed error can occur if the download process fails due to network connectivity issues.

The following parameters are arguments are supported:

- Command URI—URI to fetch the configuration file from or upload the file to (for supported URI format please refer to the CLI command “configuration fetch” for more details)
- Config file name—filename to save the configuration file to or to upload to remote location

The following commands are supported:

- BinarySwitchTo—replaces the configuration file with a new binary configuration file. This option fetches the configuration file from the URI provided in the mellanoxConfigDBCmdUri and switches to that configuration file. This command should be preceded by a reload command in order for the new configuration to apply.
- TextApply—fetches a configuration file in human-readable format and applies its configuration upon the current configuration.
- BinaryUpload—uploads a binary format configuration file of the current running configuration or an existing configuration file on the switch to the URI in the mellanoxConfigDBCmdUri command. The filename parameter indicates what configuration file on the switch to upload.
- TextUpload—uploads a human-readable configuration file of the current running configuration of an existing configuration file on the switch to the URI in the mellanoxConfigDBCmdUri command. The filename parameter indicates what configuration file on the switch to upload (same as the CLI command configuration text generate file <filename> upload).
- ConfigWrite—saves active configuration to a filename on the switch as given in the filename parameter. In case filename is “active”, active configuration is saved to the current saved configuration (same as the CLI command configuration write).
- BinaryDelete—deletes a binary based configuration file
- TextDelete—deletes a text based configuration file

6.1.8.6 Upgrading OS Software with SNMP

The OS supports upgrading its software using an SNMP SET request as described in MELLANOX-SW-UPDATE MIB.

The software upgrade command is issued via the OID mellanoxSWUpdateCmdExecute. The following options are available:

- Update—fetches the image from a specified URI (equivalent to the command “image fetch” followed by “image install”)
  The image to update from is defined by the OID mellanoxSWUpdateCmdUri. The restrictions on the URI are identical to what is supported in the CLI command “image fetch”.
- Set-Next—changes the image for the next boot equivalent to the CLI command “image boot”)
  The partition from which to boot is defined by the OID mellanoxSWUpdateCmdSetNext. The parameters for this OID are as follows:
  - 0—no change
  - 1—partition 1
  - 2—partition 2
  - 3—next partition (default)

Using the OIDs mellanoxSWUpdateCmdStatus and mellanoxSWUpdateCmdStatusString, you may view the status of the latest operation performed from the aforementioned in either integer values, or human-readable forms, respectively. The integer values presented may be as follows:

- 0—no operation
- 1-100—progress in percentage
- 101—success
- 200—failure

6.1.8.7 IF-MIB and Interface Information

The OS supports displaying information of switch ports, LAG ports, MLAG ports and VLAN interfaces on all systems via SNMP interface. This feature is enabled by default. The interface information is available in the ifTables, ifXTable and mellanoxIfVPITable.

Additionally, traps for interface up/down, and internal link suboptimal speed are enabled. It is possible to enable one or both of these traps.

Interface up/down traps are sent whenever there is a change in the interface’s operational state. These traps are suppressed for internal links when the internal link’s speed does not match the configured speed of the link (mismatch condition).

6.1.8.8 IGMP-STD-MIB Information

The system exposes IGMP snooping information via the IGMP-STD-MIB. This MIB displays IGMP snooping information only and shows a minimal view. Below are tables that represent mapping IGMP snooping information into RFC2933 MIB tables. Both tables are read-only.

igmpInterfaceTable mapping information:

Column	Mapped IGMP Snooping Information
igmpInterfaceIfIndex	VLAN interfaces indices on which snooping is enabled. This table displays VLAN interfaces only. For those VLANs that have no VLAN L3 interfaces, this table shows VLAN L2 indices.
igmpInterfaceStatus	Always Active.
igmpInterfaceQueryInterval	The IGMP snooping “query interval” field for this particular VLAN.
igmpInterfaceVersion	The IGMP snooping “version” field for this particular VLAN.
igmpInterfaceQuerier	The IGMP snooping "elected querier" field for this particular VLAN.
igmpInterfaceRobustness	The IGMP snooping "Robustness" field for this particular VLAN.
igmpInterfaceQueryMaxResponseTime	The IGMP snooping "Response interval" field for this particular VLAN.
igmpInterfaceEntry	N/A. Always zero.
igmpInterfaceQuerierExpiryTime	N/A. Always zero.
igmpInterfaceVersion1QuerierTimer	N/A. Always zero.
igmpInterfaceWrongVersionQueries	N/A. Always zero.
igmpInterfaceJoins	N/A. Always zero.
igmpInterfaceProxyIfIndex	N/A. Always zero.
igmpInterfaceGroups	IGMP snooping groups count.
igmpInterfaceLastMembQueryIntvl	N/A. Always zero.

igmpCacheTable mapping information:

Column	Mapped IGMP Snooping Information
igmpCacheAddress	Multicast group addresses registered in IGMP snooping module.
igmpCacheIfIndex	VLAN interfaces indices (L3 and L2) and regular interface indices on which this particular multicast group is registered.
igmpCacheSelf	N/A. Always “false”.
igmpCacheLastReporter	N/A. Always “0.0.0.0”.
igmpCacheUpTime	N/A. Always zero.
igmpCacheExpiryTime	N/A. Always zero.
igmpCacheStatus	Always Active.
igmpCacheVersion1HostTimer	N/A. Always zero.

6.1.9 Additional Readings and Use Cases

For more information about this feature and its potential applications, please refer to the following community posts:

- Getting Started with SNMP MIBs
- HowTo Use SNMP SET

6.2 JSON API

JavaScript Object Notation (JSON) is a machine-to-machine data-interchange format which is supported in NVIDIA Onyx CLI.

The JSON API allows executing CLI commands and receiving outputs in JSON format which can be easily parsed by the calling software.

6.2.1 Authentication

The JSON API protocol runs over HTTP/HTTPS and uses the existing web authentication mechanism.

In order to access the system via HTTP/HTTPS, an HTTP/HTTPS client is needed to send POST requests to the system.

> HTTPS access to the web-based management console needs to be enabled using the command “web https enable” to allow POST requests.

The HTTPS client must first be authenticated by sending a POST request to the following URL:

```
https://<ip-address>/admin/launch?script=rh&template=json-request&action=json-login
```

The POST request content should contain the following data (may also be saved as a file) in a JSON format:

```
{
 "username": "<user name>",
 "password": "<user password>"
}
```

After a successful login, a session ID (cookie) is returned to be used for other HTTPS requests in the system.

6.2.1.1 Authentication Example

Before sending JSON HTTPS request, the user must first authenticate.

Create a JSON format file that contains the relevant login credentials. For example, add this content to a file called "post.json":

```
{
 "username": "admin",
 "password": "admin"
}
```

Run the following from your server's shell to create a login session ID in the file: cookiejar.

```
curl -L -X POST -d @post.json -c cookiejar "http://<ip-address>/admin/launch?script=rh&template=json-request&action=json-login"
```

Upon a successful login, you will receive a reply similar to the following:

```
{
 "status": "OK",
 "status_message": "Successfully logged-in"
}
```

The session ID can now be used in all other JSON HTTPS requests to the system.

If authentication fails, the following message is received:

```
{
 "status": "ERROR",
 "status_message": "<Invalid username or password | Please provide username and password>"
}
```

You may also log in and execute commands in the same JSON request. In this case, the JSON file must be in the following format:

```
{
 "username": "<user name>",
 "password": "<user password>",
 "commands | cmd": ["<cli command 1>", "<cli command 2>"] | "<cli command>",
 "execution_type": "sync | async"
}
```

For example:

```
{
 "username": "admin",
 "password": "admin",
 "cmd":  "show fan"
}
```

If login is successful, the JSON API response appears. Otherwise, login failure response is presented.

6.2.1.2 Changing Initial Password Through JSON API

This section provides support for changing the default password through JSON API.

Expected Input

- To change the initial password, the payload will be as follows:

```
{
"username": "admin",
"password": "admin",
"initial_admin_password": "admin",
"initial_monitor_password": "monitor"
}
```

Expected Outputs

- Admin and Monitor passwords cannot be changed because they have already been changed:

```
{
"status": "ERROR",
"status_message": " 'admin' password was already set & 'monitor' password was already set"
}
```

- Admin and Monitor passwords were changed successfully:

```
{
"status_message": " <'admin' password was updated successfully> & <'monitor' password was updated
successfully> "
}
```

- Admin and Monitor passwords were not updated:

```
{
"status": "OK",
"status_message": "'admin' password was updated successfully & 'monitor' password was updated successfully"
}
```

- One of the passwords of either Admin or Monitor was changed, while the other remained the same:

```
{
"status": "<ERROR|OK>",
"status_message": " < Initial password for the 'admin' password was already set | 'admin' password was
updated successfully> "
}
```

- When the payload does not have initial passwords, check change-password nodes to see if there is no updated password return in this JSON payload:

```
{
"status": "ERROR",
"status_message": "Please set the default password for 'admin' account by using initial password
parameters"
}
```

When there is no issue with the login, flow will proceed without needing this step.

6.2.1.3 JSON API Logout

To logout, do the following:

1. Performs a POST operation on URL (the request should contain the session cookie):

```
[switch_ip]/script=rh&template=json-request&action=json-logout
```

2. The switch will remove the session and return the following JSON in the response text (in case of error, content will be relevant to the error):

```
{
"status": "OK",
```

```
"status_message": "Successfully logged-out"
}
```

3. Make sure there is no cookie. A request with an invalid cookie will respond that the cookie is invalid.

Logout Example

To logout, use the “curl” tool.

```
curl -b cookiejar "http://[switch-ip]/admin/launch?script=rh&template=json-request&action=json-logout
```

6.2.2 Sending the Request

After successful authentication, the HTTPS client can start sending JSON requests. All requests (POST and GET) should be sent to the following URL:

After the request is handled in the system the HTTPS client receives a JSON response with an indication of the request execution result. If there is data resulting from the request, it is returned as part of the response.

See “JSON Request Format” for the CLI request format.

See “JSON Response Format” for the reply format.JSON requests may also be sent using the WebUI. For more information on using the WebUI with JSON, please refer to “JSON Request Using WebUI”.

6.2.3 JSON Request Format

6.2.3.1 JSON Execution Requests

JSON execution requests are HTTPS POST requests that contain CLI commands to be executed in the system.

Execution request can contain a single command or multiple commands to be executed.

Single command execution request format:

```
{
    "cmd": "<CLI command to execute>"
}
```

Example:

```
{
    "cmd": "show interfaces ethernet 1/1"
}
```

Multiple command execution request format:

```
{
    "commands":["<CLI cmd 1>", "<CLI cmd 2>", … , <CLI cmd n>]
}
```

Example:

```
{
     "commands":
     [
          "show interfaces ethernet 1/1",
          "show interfaces ethernet 1/2"
     ]
}
```

In case of a multiple command request, the execution of the commands is done in the order they appear in the execution list. Note that the execution of a multiple command request will be stopped upon first failure. That is, in case the execution of one of the commands fails, none of the remaining commands will be executed.

6.2.3.1.1 Execution Types

Execution requests can be either synchronous (default) or asynchronous.

Synchronous requests will wait for a JSON response from the system. The synchronous request has a defined wait time after which the user will receive a timeout response. The timeout for a synchronous request is configurable by the user and is 30 seconds by default (see the CLI command "json-gw synchronous-request-timeout").

Asynchronous requests will return immediately after sending the request with a reply containing a "job_id" key. The user can use the given job ID to later query for request status and execution results. Queries for asynchronous request results are guaranteed to be accessible up to 60 seconds after the request has been completed. After the result has been successfully queried it will be deleted and will no longer be accessible (even if the result is not 60 seconds old).

To specify the execution type, the user needs to add the following key to the JSON execution request:

```
"execution_type":"<async|sync>"
```

Example:

```
{
     "execution_type":"async",
     "cmd": "show interfaces ethernet 1/1"
}
```

6.2.3.2 JSON Query Requests

JSON Query requests are HTTPS GET requests that contain a job ID parameter. Using a query request, the user can get information on the current execution state of an ongoing request or the execution results of a completed request. To send a query request, the user should add the following parameters to the JSON URL:

```
job_id=<job number>
```

Example:

```
https://<switch-ip-address>/admin/launch?script=json&job_id=<job number>
```

See "JSON Examples" for more examples.

6.2.4 JSON Response Format

> Set commands normally do not return any data or output. If a set command does return an output, it will be displayed in the “status_message” field.

6.2.4.1 Single Command Response Format

The HTTPS POST response format structure is a JSON object consisting of 4 name-value pairs as follows:

```
{
    "executed_command": "<CLI command that was executed>",
    "status" = "<OK|ERROR>",
    "status_message" = "<information on the status received>",
    "data" = {the information that was asked for in the request}
}
```

- executed_command—the CLI command that was executed in the request
- status—the result of the request execution:
  - “OK” if the execution is successful
  - “ERROR” in case of a problem with the execution
- The value type of this key is “string”.
- data—a JSON object containing the information requested. Returns an empty string if there is no data.
- status message—additional information on the received status. May be empty. The value type of this key is “string”.

Example:

```
{
    “executed_command”: “show interfaces ethernet 1/1
    "status": "OK",
    "status_message": "",
    "data":
    {
        "speed": "40GbE",
        "admin_state": "up"
    }
}
```

See “JSON Examples” for more examples.

6.2.4.2 Multiple Command Response Format

The HTTPS response format structure is a JSON object consisting of a list of JSON results. Each JSON structure in the list is structured the same as in the single command execution response (see the previous section).

However, the status field can contain in this case an additional value, “ABORTED”, in case a previous command failed. This status value indicates that the command has not been executed at all in the system.

```
{
    "results": [
    {
        "executed_command": "<…>",
        "status": "<OK|ERROR|ABORTED>",
```

```
        "status_message": "<…>",
        "data": {…}
      },
      {
        "executed_command": "<…>",
        "status": "<OK|ERROR|ABORTED>",
        "status_message": "<…>",
        "data": {…}
      },
…
      {
        "executed_command": "<…>",
        "status": "<OK|ERROR|ABORTED>",
        "status_message": "<…>",
        "data": {…}
      }
            ]
}
```

Example:

```
{
    "results": [
    {
        "executed_command": "show interfaces ethernet 1/1",
        "status": "OK",
        "status_message": ""
        "data": {"speed":"40GbE", "admin_state":"up"}
    },
    {
        "executed_command": "show interfaces ethernet 1/100",
        "status": "ERROR",
        "status_message": "wrong interfaces name",
        "data": ""
    },
    {
        "executed_command": "show interfaces ethernet 1/2",
        "status": "ABORTED",
        "status_message": "",
        "data": ""
    }
            ]
}
```

See “JSON Examples” for more examples.

6.2.4.3 Query Response Format

Response to a query request can be of two types. In case the request completes its execution, the response will be similar to the single/multiple command response format, depending on the format of the request, and will display the execution results.
In case the execution is not complete yet, the response format will be similar to the single command response format. However, the status field will contain in this case the value “PENDING” to indicate that the request is still in progress. In addition, the “executed_command” field will contain the current request command being handled by the system.

Example:

```
{
    "executed_command": "show interfaces ethernet 1/1",
    "status": "PENDING",
    "status_message": "",
    "data":""
}
```

6.2.4.4 Asynchronous Response Format

Response to an asynchronous request is similar to the HTTPS response format of the single command response. However, an additional unique field will be added, “job_id”, containing the job id number for querying the request later. The value of the job_id key is of type string.

Another difference is that the “executed_command” field will be empty.

Example:

```
{
    "executed_command": ""
    "status": "OK"
    "status_message": ""
    "data": ""
    "job_id": "2754930426"
}
```

6.2.5 Supported Commands

- Show commands
- Set commands—all non-interactive CLI set commands are supported

> Interactive commands are commands which require user interaction to complete (e.g. type “yes” to confirm). These commands are not supported by the JSON API.

6.2.6 JSON Examples

The following examples use curl (a common tool in Linux systems) to send HTTPS POST requests to the system.

6.2.6.1 Synchronous Execution Request Example

6.2.6.1.1 Single Command

This example sends a request to query the system profile.

Request (save it to a file named req.json):

```
{"cmd": "show system profile"}
```

Send the request:

```
curl -b /tmp/cookie -X POST -d @req.json "https://10.10.10.10/admin/launch?script=json"
```

When the system finishes processing the request, the user will receive a response similar to the following:

```
{
    "status": "OK",
    "executed_command": "show system profile",
    "status_message": "",
    "data": {
        "Profile": "eth",
        "Adaptive Routing": "yes",
        "Number of SWIDs": "1"
    }
}
```

6.2.6.1.2 Multiple Commands

This example sends a request to change an interface description and then queries for its status.

Request (save it to a file named req.json):

```
{"commands": ["interfaces ethernet 1/1 description test description",
"show interfaces ethernet 1/1 status"]}
```

Send the request:

```
curl -b /tmp/cookie -X POST -d @req.json "https://10.10.10.10/admin/launch?script=json"
```

When the system finishes processing the request, the user will receive a response similar to the following:

```
{
    "results": [
        {
            "status": "OK",
            "executed_command": "interfaces ethernet 1/1 description test description",
            "status_message": "",
            "data": ""
        },
        {
            "status": "OK",
            "executed_command": "show interfaces ethernet 1/1 status",
            "status_message": "",
            "data": {
                "ETH1/1": [
                    {
                        "Negotiation": "Auto",
                        "Operational state" : "Down"
                        "Speed": "Unknown",

                    }
                ]
            }
        }
    ]
}
```

6.2.6.2 Asynchronous Execution Request Example

This example sends an asynchronous request to change an interface description and then queries for its status.

Request (save it to a file named req.json):

```
{"execution_type":"async",
"commands": ["interfaces ethernet 1/1 description test description",
"show interfaces ethernet 1/1 status"]}
```

Send the request:

```
curl -b /tmp/cookie -X POST -d @req.json "https://10.10.10.10/admin/launch?script=json"
```

The system immediately returns a response similar to the following:

```
{
        "executed_command": "",
        "status": "OK",
        "status_message": "",
        "data": "",
        "job_id": "91329386"
}
```

6.2.6.3 Query Request Example

This example sends a request to query for a job ID received from a previous execution request.

The request is a an HTTPS GET operation to the JSON URL with the “job_id” parameter.

Send the request:

```
curl -b /tmp/cookie -X GET "https://10.10.10.10/admin/launch?script=json&job_id=91329386"
```

If the system is still processing the request, the user receives a response similar to the following:

```
{
        "executed_command": " interfaces ethernet 1/1 description test description ",
        "status": "PENDING",
        "status_message": "",
        "data": ""
}
```

If the system is done processing the request, the user receives a response similar to the following:

```
{
    "results": [
        {
            "status": "OK",
            "executed_command": "interfaces ethernet 1/1 description test description",
            "status_message": "",
            "data": ""
        },
        {
            "status": "OK",
            "executed_command": "show interfaces ethernet 1/1 status",
            "status_message": "",
            "data": {
                "ETH1/1": [
                    {
                        "Negotiation": "Auto",
                        "Operational state" : "Down"
                        "Speed": "Unknown",
                    }
                ]
            }
        }
    ]
}
```

6.2.6.4 Error Response Example

6.2.6.4.1 General Error

This example sends a request with an illegal JSON structure.

Request—without closing bracket “]” (save it to a file named req.json):

```
{"commands": ["interfaces ethernet 1/1 description test description",
"show interfaces ethernet 1/1 status"}
```

Send the request:

```
curl -b /tmp/cookie -X POST -d @req.json "https://10.10.10.10/admin/launch?script=json"
```

Error response:

```
{
    "status": "ERROR",
    "executed_command": "",
    "status_message": "Handle request failed. Reason:\nIllegal JSON structure found in given JSON data.
\nExpecting , delimiter: line 1 column 95 (char 94)",
    "data": ""
}
```

6.2.6.4.2 Multiple Command Request Failure

This example sends a multiple command request where one command fails.

Request—with a non-existing interface (1/200) (save it to a file named req.json):

```
{
        "execution_type": "sync",
        "commands": [   "interfaces ethernet 1/1 speed 25.0 Gbps",
                        "interfaces ethernet 1/200 speed 25.0 Gbps",
                        "interfaces ethernet 1/3 speed 25.0 Gbps"]
}
```

Send the request:

```
curl -b /tmp/cookie -X POST -d @req.json "https://10.10.10.10/admin/launch?script=json"
```

Error response:

```
{
    "results": [
        {
            "status": "OK",
            "executed_command": "interfaces ethernet 1/1 speed 25.0 Gbps ",
            "status_message": "",
            "data": ""
        },
        {
            "status": "ERROR",
            "executed_command": "interfaces ethernet 1/200 speed 25.0 Gbps",
            "status_message": "% 1st Interface does not exist",
            "data": ""
        },
        {
            "status": "ABORTED",
            "executed_command": "interfaces ethernet 1/3 speed 25.0 Gbps",
            "status_message": "",
            "data": ""
        }
    ]
}
```

6.2.7 JSON Request Using WebUI

The
NVIDIA Onyx

WebUI also allows users to send JSON HTTPS POST and GET requests.

Log into the WebUI, go to the “Setup” tab, and select “JSON API” from the left side menu.

> This section is displayed only if JSON API is enabled using the command “json-gw enable”.

6.2.7.1 To Execute a JSON Request

1. Choose “Execute JSON command”.
2. Choose the “execution_type” from the drop down list.
3. In the “commands” field, type the CLI command(s) to execute.
   Use the “+” and “-” buttons to add or remove additional commands to the request.
4. Click “Submit”.

The JSON response is then shown in the “JSON Response” box below.

The HTTPS method (HTTPS POST in this instance) and the URL used to send the request will be displayed next to the “HTTPS Method” and “URL” field respectively.

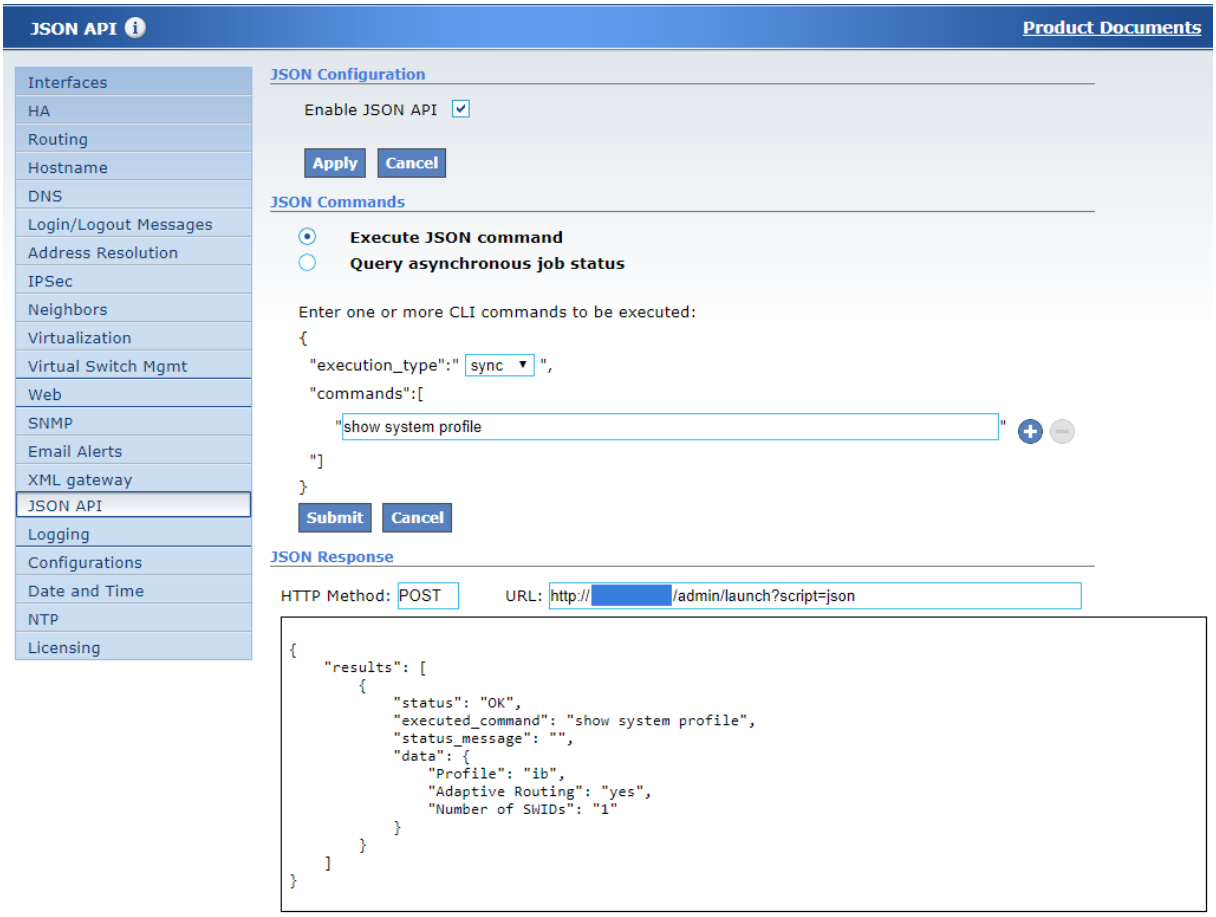

6.2.7.2 To Query an Asynchronous JSON Request

1. Choose “Query asynchronous job status”.
2. Type the job ID in the “Job ID” text box.
3. Press “Query Status”.

The JSON response is then shown in the “JSON Response” box below.
The HTTPS method (HTTPS GET in this instance) and the URL used to send the request will be displayed next to the “HTTPS Method” and “URL” field respectively.

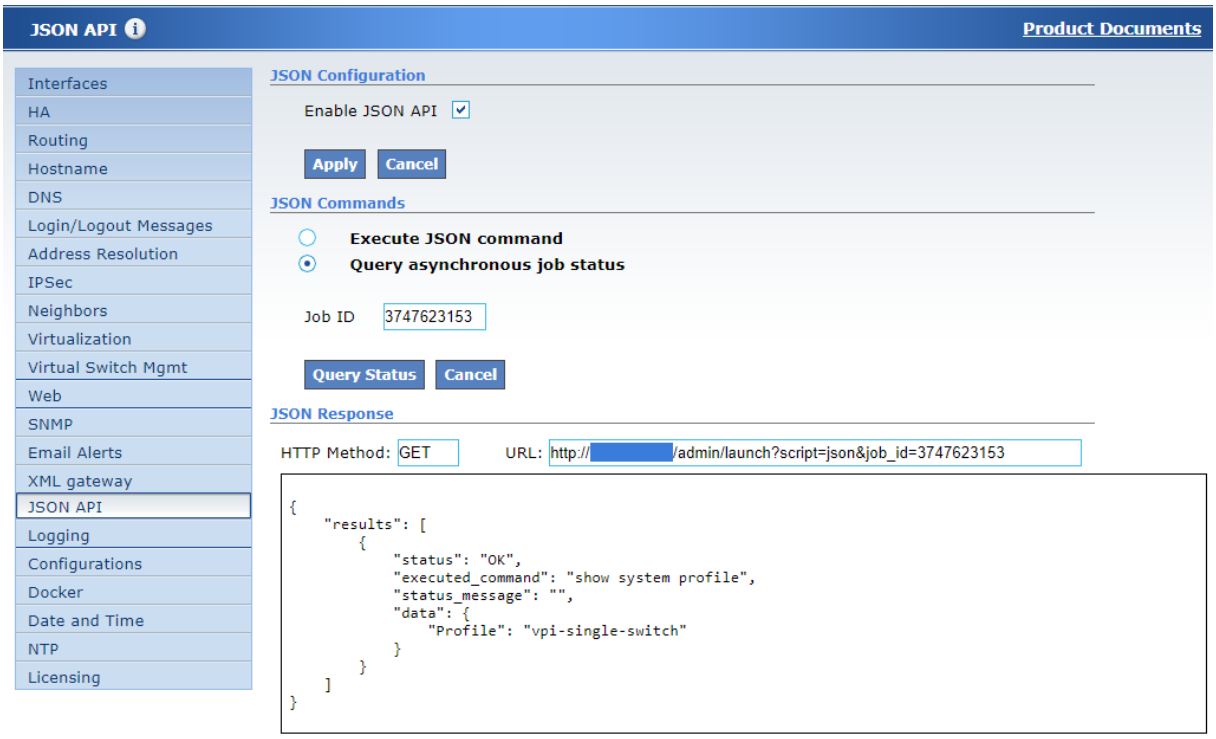

6.2.8 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community post:

- Getting Started With JSON API

6.3 Network Management Interface Commands

- 6.3.1 SNMP
  - 6.3.1.1 snmp-server auto-refresh
  - 6.3.1.2 snmp-server cache enable
  - 6.3.1.3 snmp-server community
  - 6.3.1.4 snmp-server contact
  - 6.3.1.5 snmp-server enable
  - 6.3.1.6 snmp-server engineID reset
  - 6.3.1.7 snmp-server enable mult-communities
  - 6.3.1.8 snmp-server enable notify
  - 6.3.1.9 snmp-server enable set-permission
  - 6.3.1.10 snmp-server host disable
  - 6.3.1.11 snmp-server host informs
  - 6.3.1.12 snmp-server host traps
  - 6.3.1.13 snmp-server listen
  - 6.3.1.14 snmp-server notify
  - 6.3.1.15 snmp-server port

- 6.3.1.16 snmp-server user
- 6.3.1.17 show snmp
- 6.3.1.18 show snmp auto-refresh
- 6.3.1.19 show snmp engineID
- 6.3.1.20 show snmp set-permission
- 6.3.1.21 show snmp user

- 6.3.2 JSON API
  - 6.3.2.1 json-gw enable
  - 6.3.2.2 json-gw synchronous-request-timeout
  - 6.3.2.3 show json-gw

6.3.1 SNMP

6.3.1.1 snmp-server auto-refresh

	snmp-server auto-refresh {enable | interval <time>} no snmp-server auto-refresh enable Configures SNMPD refresh settings. The no form of the command disables SNMPD refresh mechanism.	
Syntax Description	enable	Enables SNMPD refresh mechanism.
	interval	Sets SNMPD refresh interval.
	time	Range: 20-500 seconds
Default	Enabled Interval—60 seconds	
Configuration Mode	config	
History	3.2.3000 3.4.1100: Added “time” parameter and updated notes	
Example	`switch (config) # snmp-server auto-refresh interval 120`	
Related Commands	show snmp	
Notes	• When configuring an interval lower than 60 seconds, the following warning message appears asking for confirmation: “Warning: this configuration may increase CPU utilization, Type 'YES' to confirm: YES • When disabling SNMP auto-refresh, information is retrieved no more than once every 60 seconds just like SNMP tables that do not have an auto-refresh mechanism	

6.3.1.2 snmp-server cache enable

	snmp-server cache enable no snmp-server cache enable Enables SNMP cache if auto-refresh is disabled. The no form of the command disables SNMP cache if auto-refresh is disabled.
Syntax Description	N/A
Default	Enabled

Configuration Mode	config
History	3.7.0000
Example	`switch (config) # snmp-server cache enable`
Related Commands	show snmp auto-refresh snmp-server auto-refresh enable
Notes	• If SNMP auto-refresh is enabled, the value of cache is meaningless • If SNMP cache is disabled, every SNMP request gets updated data

6.3.1.3 snmp-server community

	snmp-server community <community> [ro | rw] no snmp-server community <community> Sets a community name for either read-only or read-write SNMP requests. The no form of the command sets the community string to default.	
Syntax Description	community	Community name
	ro	Sets the read-only community string
	rw	Sets the read-write community string
Default	Read-only community: “public” Read-write community: “”	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # snmp-server community private rw`	
Related Commands	show snmp	
Notes	• If neither the “ro” or the “rw” parameters are specified, the read-only community is set as the default community • If the read-only community is specified, only queries can be performed • If the read-write community is specified, both queries and sets can be performed	

6.3.1.4 snmp-server contact

	snmp-server contact <contact-name> no snmp-server contact Sets a value for the sysContact variable in MIB-II. The no form of the command resets the parameter to its default value.	
Syntax Description	contact-name	Contact name
Default	“”	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # snmp-server contact my-name`	
Related Commands	show snmp	

Notes	

6.3.1.5 snmp-server enable

	snmp-server [vrf <vrf-name>] enable [force] no snmp-server [vrf <vrf-name>] enable Enables SNMP-related functionality (SNMP engine, and traps). The no form of the command disables the SNMP server.
Syntax Description	vrf name—Describes VRF name for snmp-server. If "vrf" parameter is not specified, the "default" VRF will be used force—Restarts SNMP server with previous VRF context even if it was already enabled in using other VRF.
Default	SNMP is enabled by default
Configuration Mode	config
History	 3.1.0000 3.9.2000—Added VRF option
Example	`switch (config) # snmp-server enable`
Related Commands	show snmp
Notes	SNMP server can be enabled only in one VRF at a time.

6.3.1.6 snmp-server engineID reset

	snmp-server engineID reset Resets the SNMPv3 engine ID to be node unique.
Syntax Description	N/A
Default	Default engineID is unchanged
Configuration Mode	config
History	 3.6.6102
Example	`switch (config) # snmp-server engienID reset`
Related Commands	show snmp engineID
Notes	Changing system profile or performing “reset factory...” causes the engine ID to change to the new node-unique one.

6.3.1.7 snmp-server enable mult-communities

	snmp-server enable mult-communities no snmp-server enable mult-communities Enables multiple communities to be configured. The no form of the command disables multiple communities to be configured.
Syntax Description	N/A
Default	SNMP server multi-communities are disabled by default

Configuration Mode	config
History	3.1.0000
Example	`switch (config) # snmp-server enable mult-communities`
Related Commands	show snmp
Notes	

6.3.1.8 snmp-server enable notify

	snmp-server enable notify no snmp-server enable notify Enables sending of SNMP traps and informs from this system. The no form of the command disables sending of SNMP traps and informs from this system.
Syntax Description	N/A
Default	SNMP notifies are enabled by default
Configuration Mode	config
History	3.1.0000
Example	`switch (config) # snmp-server enable notify`
Related Commands	show snmp
Notes	SNMP traps are only sent if there are trap sinks configured with the “snmp-server host...” command, and if these trap sinks are themselves enabled.

6.3.1.9 snmp-server enable set-permission

	snmp-server enable set-permission <MIB-name> no snmp-server enable set-permission <MIB-name> Allows SNMP SET requests for items in a specified MIB. The no form of the command disallows SNMP SET requests for items in a specified MIB.
Syntax Description	N/A
Default	SNMP MIBs are all given permission for SET requests by default
Configuration Mode	config
History	3.6.3004
Example	`switch (config) # snmp-server enable set-permission MELLANOX-SW-UPDATE`
Related Commands	show snmp set-permission
Notes	

6.3.1.10 snmp-server host disable

	snmp-server host <ip-address> disable no snmp-server host <ip-address> [disable] Temporarily disables sending of all notifications to this host. The no form of the commands resumes sending of all notifications to this host.	
Syntax Description	ip-address	IPv4 or IPv6 address
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # snmp-server host 10.10.10.10 disable`	
Related Commands	show snmp snmp-server enable	
Notes		

6.3.1.11 snmp-server host informs

	snmp-server host [vrf <vrf-name>] <ip-address> informs [<community> | port <port> | version 2c | version 3 {engineID <engineID> | user <name> {auth <hash-type> <auth-password> [priv <privacy-type> [<priv-password>]] | encrypted auth ... | prompt auth ...}}] no snmp-server host <ip-address> informs port Send SNMP v2c informs to this host with the default trap community. The no form of the commands removes a host from which SNMP traps should be sent.	
Syntax Description	vrf-name—Describes the VRF name for NTP daemon. If the VRF parameter is not specified, the "default" VRF will be used implicitly.	
	IP address	IPv4 or IPv6 address.
	community	Specifies trap community string.
	port	Overrides default UDP port for this trap sink.
	version	Specifies the SNMP version of traps to send to this host.
	engineID	Specifies engine ID of this inform sink.
	user	Specifies username for this inform sink.
	auth	Configures SNMPv3 security parameters, specifying passwords in plaintext on the command line (passwords are always stored encrypted).
	hash-type	• MD5 • SHA
	auth-password	Plaintext password to use for authentication. If “priv” is not specified the default privacy algorithm is used with the same privacy password as that specified for authentication.
	priv	Specifies SNMPv3 privacy settings for this user.
	privacy-type	• aes-128—uses AES-128 encryption for privacy • des—uses DES encryption for privacy
	priv-password	Plaintext password to use for privacy. If not specified, then auth-password is used.

	encrypted	Configure SNMPv3 security parameters specifying passwords in encrypted form.
	prompt	Configure SNMPv3 security parameters specifying passwords securely in follow-up prompts rather than on the command line.
Default	community—public UDP port—162 version—3	
Configuration Mode	config	
History	 3.2.1050 3.9.2000—Added VRF option	
Example	`switch (config) # snmp-server host 1.1.1.1 informs version 3 engineID 0x800041da0464326536393265343230313 5 user test auth md5 password priv aes-128 password`	
Related Commands	show snmp snmp-server enable snmp-server host informs version 3	
Notes	Multiple snmp-hosts can be configured in different VRF`s at the same time.	

6.3.1.12 snmp-server host traps

	snmp-server host [vrf <vrf-name>] <ip-address> traps [<community> | port <port> | version {1 | 2c} | version 3 {user <name> {auth <hash-type> <auth-password> [priv <privacy-type> [<priv-password>]] | encrypted auth ... | prompt auth ...}}] no snmp-server host <ip-address> traps port Send SNMP v2c traps to this host with the default trap community. The no form of the commands removes a host from which SNMP traps should be sent.	
Syntax Description	vrf-name—Describes the VRF name for NTP daemon. If the VRF parameter is not specified, the "default" VRF will be used implicitly.	
	ip-address	IPv4 or IPv6 address.
	community	Specifies trap community string.
	port	Overrides default UDP port for this trap sink.
	version	Specifies the SNMP version of traps to send to this host.
	user	Specifies username for this inform sink.
	auth	Configures SNMPv3 security parameters, specifying passwords in plaintext on the command line (passwords are always stored encrypted).
	hash-type	• MD5 • SHA
	auth-password	Plaintext password to use for authentication. If “priv” is not specified the default privacy algorithm is used with the same privacy password as that specified for authentication.
	priv	Specifies SNMPv3 privacy settings for this user.
	privacy-type	• aes-128—uses AES-128 encryption for privacy • des—uses DES encryption for privacy
	priv-password	Plaintext password to use for privacy. If not specified, then auth-password is used.

	encrypted	Configure SNMPv3 security parameters, specifying passwords in encrypted form.
	prompt	Configure SNMPv3 security parameters, specifying passwords securely in follow-up prompts, rather than on the command line.
	vrf-name—Describes VRF name for snmp-server. If "vrf" parameter is not specified, the "default" VRF will be used	
Default	community—public UDP port—162 version—3	
Configuration Mode	config	
History	3.1.0000 3.9.2000—Added VRF option	
Example	`switch (config) # snmp-server host 1.1.1.1 informs version 3 user test auth md5 password priv aes-128 password`	
Related Commands	show snmp snmp-server enable snmp-server host informs version 3	
Notes	Multiple snmp-hosts can be configured in different VRF`s at the same time.	

6.3.1.13 snmp-server listen

	snmp-server listen {enable | interface <ifName>} no snmp-server listen {enable | interface <ifName>} Configures SNMP server interface access restrictions. The no form of the command disables the listen interface restricted list for SNMP server.	
Syntax Description	enable	Enables SNMP interface restrictions on access to this system
	ifName	Adds an interface to the “listen” list for SNMP server. For example: “mgmt0”, “mgmt1”
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # snmp listen enable`	
Related Commands	show snmp	
Notes	If enabled, and if at least one of the interfaces listed is eligible to be a listen interface, then SNMP requests will only be accepted on those interfaces. Otherwise, SNMP requests are accepted on any interface.	

6.3.1.14 snmp-server notify

	snmp-server notify {community <community> | event <event name> | port <port> | send-test} no snmp-server notify {community | event <event name> | port} Configures SNMP notifications (traps and informs). The no form of the commands negate the SNMP notifications.	
Syntax Description	community	Sets the default community for traps sent to hosts which do not have a custom community string set
	event	Specifies which events will be sent as traps
	port	Sets the default port to which traps are sent
	send-test	Sends a test trap
Default	All informs and traps are enabled community—public UDP port—162	
Configuration Mode	config	
History	3.1.0000 3.2.1050: Changed traps to notify	
Example	`switch (config) # snmp-server community public`	
Related Commands	show snmp show snmp events	
Notes	• This setting is only meaningful if traps are enabled, though the list of hosts may still be edited if traps are disabled • Refer to Mellanox MIB file for the list of supported traps	

6.3.1.15 snmp-server port

	snmp-server port <port> no snmp-server port Sets the UDP listening port for the SNMP agent. The no form of the command resets the parameter to its default value.	
Syntax Description	port	UDP port
Default	161	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # snmp-server port 1000`	
Related Commands	show snmp	
Notes		

6.3.1.16 snmp-server user

	snmp-server user {admin | <username>} v3 {[encrypted] auth <hash-type> <password> [priv <privacy-type> [<password>]] | capability <cap> | enable <sets> | prompt auth <hash-type> [priv <privacy-type>] | require-privacy} no snmp-server user {admin | <username> } v3 {[encrypted] auth <hash-type> <password> [priv <privacy-type> [<password>]] | capability <cap> | enable <sets> | prompt auth <hash-type> [priv <privacy-type>]} Specifies an existing username, or a new one to be added. The no form of the command disables access via SNMP v3 for the specified user.	
Syntax Description	v3	Configures SNMPv3 users.
	auth	Configures SNMPv3 security parameters, specifying passwords in plaintext on the command line (note: passwords are always stored encrypted). Available hash-type options are: <md5|sha|sha224|sha256|sha384|sha512>.
	capability	Sets capability level for SET requests.
	enable	Enables SNMPv3 access for this user.
	encrypted	Configures SNMPv3 security parameters, specifying passwords in encrypted form.
	prompt	Configures SNMPv3 security parameters, specifying passwords securely in follow-up prompts, rather than on the command line.
	require-privacy	Requires privacy (encryption) for requests from this user.
	priv	Configures SNMPv3 security parameters, specifying which protocol to use for traffic encryption. Available priv-type options: <des|3des|aes-128|aes-192|aes-256>.
Default	No SNMP v3 users defined	
Configuration Mode	config	
History	3.1.0000 3.7.0000 3.8.1000: Syntax updated	
Example	`switch (config) # snmp-server user admin v3 enable`	
Related Commands	show snmp user	

Notes	• The username chosen here may be anything that is valid as a local UNIX username (alphanumeric, plus '-', '_', and '.'), but these usernames are unrelated to, and independent of, local user accounts. That is, they need not have the same capability level as a local user account of the same name. Note that these usernames should not be longer than 31 characters, or they will not work. • The hash algorithm specified is used both to create digests of the authentication and privacy passwords for storage in configuration, and also in HMAC form for the authentication protocol itself • There are three variants of the command, which branch out after the “v3” keyword. If “auth” is used next, the passwords are specified in plaintext on the command line. If “encrypted” is used next, the passwords are specified encrypted (hashed) on the command line. If “prompt-pass” is used, the passwords are not specified on the command line the user is prompted for them when the command is executing. If “priv” is not specified, only the auth password is prompted for. If “priv” is specified, the privacy password is prompted for; entering an empty string for this prompt will result in using the same password specified for authentication. • AES privacy type encryption using the newest algorithm, which means we use aes-blumenthal. For more information see http://www.snmp.com/eso/esoConsortiumMIB.txt. • No more than 30 SNMPv3 users are allowed in the database

6.3.1.17 show snmp

	show snmp [events | host] Displays SNMP-server configuration and status.	
Syntax Description	events	SNMP events
	host	List of notification sinks
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000 3.6.8008—Updated example 3.9.2000—Updated example, adding VRF option	

Example	switch (config) # show snmp SNMP enabled : no SNMP port : 161 System contact : Test System location: Boston VRF name : mgmt Read-only communities: public Read-write communities: good Interface listen enabled: yes Listen Interfaces: Interface: mgmt0 switch (config) # show snmp host Notifications enabled : yes Default notification community: public Default notification port : 162 Notification sinks: 20.20.20.20: Enabled : yes Port : 162 (default) Notification type: SNMP v2c trap Community : public (default) VRF : other 10.10.10.10: Enabled : yes Port : 162 (default) Notification type: SNMP v2c inform Community : public (default) VRF : default
Related Commands	show snmp
Notes	

6.3.1.18 show snmp auto-refresh

	show snmp auto-refresh Displays SNMPD refresh mechanism status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.6.6000: Updated example 3.7.0000: Updated example
Example	switch (config) # show snmp auto-refresh SNMP auto refresh: Auto-refresh enabled: yes Refresh interval (sec): 60 Cache enabled: yes Auto-Refreshed tables: ifTable ifXTable mellanoxIfVPITable
Related Commands	snmp-server auto-refresh

Notes	

6.3.1.19 show snmp engineID

	show snmp engineID Displays SNMPv3 engine ID key.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.6102
Example	switch (config) # show snmp engineID Local SNMP engineID: 0x80004f4db1dd435e80accf4a4d4d3031
Related Commands	snmp-server engineID
Notes	

6.3.1.20 show snmp set-permission

	show snmp set-permission Displays SNMP SET permission settings.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.3004
Example	switch (config) # show snmp set-permission --- MIB Name Set Enable --- MELLANOX-CONFIG-DB-MIB yes MELLANOX-EFM-MIB yes MELLANOX-POWER-CYCLE yes MELLANOX-SW-UPDATE no RFC1213-MIB no
Related Commands	snmp-server enable set-permission
Notes	

6.3.1.21 show snmp user

	show snmp user Displays SNMP user information.
Syntax Description	N/A
Default	N/A
Configuration Mode	config

History	3.1.0000 3.6.8008: Updated example

Example	`switch (config) # show snmp user`
	`User name: Hendrix`
	`  Enabled overall:  yes`
	`  Authentication type:  sha`
	`  Privacy type:  des`
	`  Authentication password:  (set)`
	`  Privacy password:  (set)`
	`  Require privacy: yes`
	`  SET access:`
	`  Enabled:  yes`
	`  Capability level:  admin`

Related Commands	show snmp
Notes	

6.3.2 JSON API

6.3.2.1 json-gw enable

	json-gw enable no json-gw enable Enables the JSON API. The no form of the command disables the JSON API.
Syntax Description	N/A
Default	JSON API is enabled
Configuration Mode	config
History	3.6.3004
Example	`switch (config) # json-gw enable`
Related Commands	show json-gw
Notes	

6.3.2.2 json-gw synchronous-request-timeout

	json-gw synchronous-request-timeout <timeout-value> no json-gw synchronous-request-timeout Defines a timeout value for synchronous JSON requests (in seconds). The no form of the command returns the timeout value to its default.	
Syntax Description	timeout-value	Define a timeout value for synchronous JSON requests Range: 0-4294967295
Default	JSON API is enabled	
Configuration Mode	config	
History	3.6.3004	
Example	`switch (config) # json-gw synchronous-request-timeout 100`	

Related Commands	show json-gw
Notes	

6.3.2.3 show json-gw

	show json-gw Displays the JSON API setting.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	 3.6.3004 3.6.4000: Updated example

Example	`switch (config) # show json-gw`
	
	`JSON Gateway enabled: yes`
	`Synchronous request timeout: 30`
	`JSON API version: 1.0`

Related Commands	json-gw enable json-gw synchronous-request-timeout <time out value>
Notes	

7 Virtualization

NVIDIA Onyx allows the user to run their own applications on a Linux docker image embedded in the switch software. The container is a pure application sandbox with resource isolation of both memory and compute from the system code/NOS.

Docker container implementation in the OS enhances its VM support to provide a new set of capabilities:

- Network traffic access
  Docker containers are implemented in the OS in the same name-space as the network devices allowing the software to send and receive packets from the switch ports by opening a standard Linux socket over the network devices and using an IP address assigned to the device via the legacy management interface (e.g., JSON over HTTP).

  > It is recommended to assign a unique port number to the Linux socket to prevent ambiguity of applications between the container and the OS.

- Calling the SDK interfaces
  Applications running in the docker container are able to implement a set of tools pertaining only to the container such as telemetry features within the network devices. By calling the switch SDK APIs, it can also read data that is not exposed in the OS user interface, or register to receive events that occur in the system (e.g., port up/down).

  > The container implementation does not limit the container developer from calling the SDK to set parameters. However this is strongly discouraged as it may cause unexpected system behavior where the OS and the container application manage the same resources.

- Query the Linux tables provisioned by OS such as neighbor cache, routing tables, L3 interfaces attributes etc.

7.1 Limiting the Container's Resources

It is possible to configure multiple containers in dockers, however, they would compete for the same memory and compute resources allocated by the switch software (varies for different systems). To ensure system stability and that no random process is killed to free up memory, it is strongly recommended that all resource configurations done in the container utilize OS user interfaces such as JSON/SNMP and take advantage of the internal loopback interface.

7.1.1 Memory Resources Allocation Protocol

The Linux docker supports a hard limit to control memory resource allocation which limits the container to a given amount of user/system memory.

To set the amount of memory allocated to the container, run the following command:

```
switch (config) # docker start imagename latestver containername init memory 25 label newlabel privileged sdk
network docker usb-mount
```

7.1.2 CPU Resource Allocation Protocol

Containers have unrestricted access to the host machine's CPU cycles but it is possible to set a number of constraints to limit the containers' access.

To set up limitations or regulate the containers access to CPU resources, run the following command:

```
docker start imagename latestver containername init cpus 0.2 label new_label privileged sdk network
```

7.2 Upgrade Ramifications

7.2.1 Changing Docker Storage Driver

As a result of the upgrade, the docker's storage driver changes, which may cause a few additional changes:

- The containers and docker images become inaccessible to the user (the docker process will not run)
- The user can reach their old containers after a rollback procedure
- The "no docker" command erases all containers and images, including those that were reachable after rollback. Rollbacking after running the "no docker" command may result in failure to create configured containers from unknown images.
- The user is advised to execute the "no docker" command at some point in order to clear unused disk space
- It is possible to reload the Docker images after upgrade with the command: docker load <image_name>_<image_version>.img.gz
- The images are presented with tab-tab after "docker load " (in cli)
- It is also possible to load the images after rollback after "no docker" was execute. That means that containers can be restarted after upgrade/rollback if their images are loaded (with "docker load").

It is possible to move containers from the current version to the updated one by executing the following steps:

Before upgrade:

1. Save the container as an image—run the command: "docker commit <container_name> <new_image_name> <new_image_version>". For example: docker commit my_name my_image my_version. You can see the new image by running: "show docker images".
2. Save the image—run the command: "docker save <image_name> <image_version> <file_name-optional>". For example: docker save my_image my_version.
3. Upload the image—save the image to a local repository by running: "image upload <image_file_name> <destination_path>". For example: image upload my_image_my_version.img.gz scp://username:password@fit150/auto/my_dir. The <image_file_name> is presented after clicking tab-tab.

After upgrade:

1. Start docker—run the "no docker shutdown" command.

2. Fetch the restored image—run the “image fetch <file_name>” command. For example: image fetch scp://username:password@fit150/auto/my_dir/my_image_my_version.img.gz
3. Load the image—run the “docker load <image_file_name>” command. For example: docker load my_image_my_version.img.gz
4. Start a container with the defined image—now that the image with all the content from the container is available in the new environment, start a container with this image. Run the command: “docker start <image_name> <image version> <docker_name> <starting_point>| privileged | label | memory | cpus | usb-mount”. For example: docker start my_image my_version new_container now

> After an upgrade operation there is a need to rerun copy-sdk command (in case in use).

7.3 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community post:

- HowTo Deploy Docker Container with DHCP Service

7.4 Docker Containers Commands

7.4.1 docker

	docker [vrf <vrf-name>] [logging-level <log-level>] [force [keep-images]] no docker [vrf <vrf-name>] Enables dockers then enters docker configuration context. The no form of the command disables dockers, removes configuration, and deletes all containers and docker images.
Syntax Description	• vrf name—Describes docker daemon VRF context, impacts fetching images and running containers. If "vrf" parameter is not specified, the "default" VRF will be used. • force—Restarts docker using past VRF context even if it was already enabled in other VRF context. • keep-images—Will not remove docker images while switching VRF context. • log-level—logging-level for docker. Possible levels: debug error, fatal info, warn
Default	N/A
Configuration Mode	config
History	3.6.2940 3.9.2000—Added VRF option 3.9.2300—Added log-level and keep-images option
Example	`switch (config) # docker`
Related Commands	

Notes	• Logging-level parameter is applicable when docker is "not-started" state or with "force" flag. If not specified, set warning level. • Only one configured instance of docker can be in the system. Moving docker between VRFs leads to restarting the docker daemon and a loss of running, cached containers and images. Pulled image can be preserved with the command "docker save".

7.4.2 docker login

	docker login <username> <cleartext password> [server <server address>] Logs in to remote docker repositories.	
Syntax Description	username	Username
	cleartext password	There are 2 options to enter password using the above command: 1. In command—cleartext 2. Using interactive shell—entering all needed input except the password will prompt the user to provide a password which will not be visible while typing. (masked by *)
	server	The "server" field is not mandatory. In case it is not present, the docker will try to login into docker hub repository.
Default	N/A	
Configuration Mode	config	
History	3.9.1600	
Example	`switch (config) # docker login abcd 1234`	
Related Commands	show docker login	
Notes		

7.4.3 docker logout

	docker logout [server <server address>] Logs out from remote server.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.9.1600
Example	`switch (config) # docker logout`
Related Commands	
Notes	• There is no need to provide username as only a single user can be connected to a specific server in any given time

7.4.4 commit

	commit <container-name> <image-name> <image-version> Creates a new image from a running container.	
Syntax Description	container-name	Name of the running container to commit (limited to 180 characters)
	image-name	Name of the new image to be created
	image-version	Version of the new image to be created
Default	N/A	
Configuration Mode	config docker	
History	3.6.2940 3.6.8008: Added new character limitation for container-name	
Example	`switch (config docker) # commit mycontainer test latest`	
Related Commands		
Notes		

7.4.5 copy-sdk

	copy-sdk The command provides access to the switch SDK APIs giving applications running on docker access to the switch hardware.
Syntax Description	N/A
Default	N/A
Configuration Mode	config docker
History	3.6.4110 3.8.1000: Updated notes 3.8.2100: Updated notes
Example	`switch (config docker) # copy-sdk`
Related Commands	
Notes	• Copying SDK files to a USB mounted folder is not allowed • After an upgrade operation there is a need to rerun copy-sdk command (in case in use).

7.4.6 remove image

	remove image <image-name> <image-version> Removes an image from the Linux docker service.	
Syntax Description	image-name	Name of the new image to be deleted
	image-version	Version of the new image to be deleted
Default	N/A	

Configuration Mode	config docker
History	3.6.3520 3.6.2940
Example	`switch (config docker) # remove image test latest`
Related Commands	docker
Notes	

7.4.7 exec

	exec <container-name> <program-executable> Executes a program within a running container.	
Syntax Description	container-name	Name of the running container to commit (limited to 180 characters)
	program-executable	Linux command
Default	N/A	
Configuration Mode	config docker	
History	3.6.3520 3.6.2940	
Example	`switch (config docker) # exec mycontainer "ls -la"`	
Related Commands	docker	
Notes		

7.4.8 label

	label <label name> no label <label name> Creates a label which can be used as a shared storage between containers. The no form of the command removes the label.
Syntax Description	N/A
Default	N/A
Configuration Mode	config docker
History	3.6.4110
Example	`switch (config docker) # label new_label`
Related Commands	
Notes	

7.4.9 load

	load <image-name> Loads an image from a TAR archive.

Syntax Description	image-name	Name of the TAR image to be loaded
Default	N/A	
Configuration Mode	config docker	
History	3.6.2940	
Example	`switch (config docker) # load test`	
Related Commands	docker	
Notes		

7.4.10 pull

	pull <image-name>[:<version>] Pulls a docker image from a docker repository.	
Syntax Description	image-name	Image name Format: Name:Version If only “Name” is provided, “version” defaults to latest
Default	N/A	
Configuration Mode	config docker	
History	3.6.2940	

Example	`switch (config docker) # pull test`
	`Using default tag: latest`
	`latest: Pulling from library/test`
	`45a2e645736c: Pull complete`
	`Digest:`
	`sha256:c577af3197aacedf79c5a204cd7f493c8e07ffbce7f88f7600bf19c688c38799`
	`Status: Downloaded newer image for test:latest`	

Related Commands	docker	
Notes		

7.4.11 save

	save <image-name> <image-version> <filename> Saves an image to a TAR archive.	
Syntax Description	image-name	Image name
	image-version	Image version
	filename	Name of the file in which to save the image
Default	N/A	
Configuration Mode	config docker	
History	3.6.2940 3.6.8008: Updated command syntax	

Example	`switch (config docker) # save busybox latest my_image`
	
	`Saving and compressing image: busybox version: latest`
	`this could take a while...`
	
	`switch (config docker) #`	

Related Commands	docker docker load
Notes	After the file is created, the filename gets appended a *.gz suffix.

7.4.12 shutdown

	shutdown no shutdown Stops all docker containers, and deletes all non-auto containers. The no form of the command enables the docker Linux service and runs all configured auto-start containers
Syntax Description	N/A
Default	N/A
Configuration Mode	config docker
History	3.6.2940
Example	`switch (config docker) # no shutdown`
Related Commands	docker
Notes	

7.4.13 start

	start <image-name> <image-version> <container-name> <starting-point> [privileged {network | sdk}] [cpus <max-cpu-resources>] [memory <max-memory>] [usb-mount] [host-trust [user <username>]] [logging-facility <logging-facility-level>] [user-env <env-string>] no start <container-name> Starts a new container from an image. The no form of the command stops a running docker container.	
Syntax Description	image-name	Name of the new image to start.
	image-version	Version of the image to start.
	container-name	Name of the running container to commit (limited to 180 characters).
	privileged	• network—adds network privileges to the container (--privilege flag) • sdk—adds required mounts to use the switch SDK from the container
	starting-point	• init—persistent, start the container after boot, when system initialization is done • data-path-ready—persistent, start the container after boot, when data-path is ready to be configured • now—start the container now, this is not persistent • now-and-data-path-ready—starts the container now and after boot, when data-path is ready to be configured • now-and-init—starts the container now and after boot, when system configuration is done • ptp-ready—persistent, start the container after boot, when protocol PTP is ready to be configured

	cpus	Sets how much of the available CPU resources a container can use (e.g., “cpus 1.5” guarantees at most one and a half of the available CPUs for the container).
	memory	Sets the maximum amount of memory the container can use in MB. The minimum amount of memory to configure is 4MB.
	usb-mount	Enables USB mount to the docker container.
	host-trust	Allows SSH operation from within the container to localhost without the need to supply password.
	logging-facility-level	Available Parameters: auth, authpriv daemon, ftp, kern, local0, local1, local2, local3, local4, local5, local6, local7, lpr, mail, news, syslog, user, uucp
	env-string	Up to 16 user-defined environment variables. User-defined environment variable are separated by a comma (e.g., key1=value1,key2=value2)
Default	N/A	
Configuration Mode	config docker	
History	3.6.2940 3.6.3520: Added “privileged” parameter 3.6.8008: Added the options “now-and-data-path-ready” and “now-and-init”, new character limitation for container-name, and updated the description of the parameter “memory” 3.7.0000—Added “ptp-ready” option 3.8.1000; Updated syntax description 3.9.2000: Added host-trust option which adds support for SSH operation from within the container to localhost without the need to supply password (when activating host-trust without supplying user, user admin will be used). 3.9.2300: Added logging-facility and user-env options	

Example	`switch (config docker) # start centos latest test now`
	
	`Starting docker container. Please wait (this can take a minute)...`
	
	`switch (config) # docker start imagename latestver containername init cpus 0.2 memory 25`	

Related Commands	docker	
Notes	• The no form of the command removes the container if it is not persistent. • If trust is set and username is not given, user admin will be chosen by default.	

7.4.14 image upload

	image upload <filename> [vrf <vrf-name>] <upload_url> Uploads an image file to a remote host.	
Syntax Description	filename	Name of file
	vrf-name—Describes VRF context that should be used for this transfer. If not specified, the “default” VRF is used.	
	upload_url	FTP, TFTP, SCP and SFTP are supported (e.g., scp://username[:password]@hostname-or-ip/path/filename)
Default	N/A	

Configuration Mode	config
History	3.6.29403.9.2000—Added VRF option
Example	`switch (config) # image upload centos.img.gz scp://username:password@192.168.10.125/var/www/html/<image_name>`
Related Commands	
Notes	

7.4.15 file image upload

	file image upload <filename> [vrf <vrf-name>] <upload_url> Uploads a file to a remote host.	
Syntax Description	filename	Name of file
	vrf-name—Describes VRF context that should be used for this transfer. If not specified, the “default” VRF is used.	
	upload_url	FTP, TFTP, SCP and SFTP are supported (e.g., scp://username[:password]@hostname/path/filename)
Default	N/A	
Configuration Mode	config	
History	3.6.29403.9.2000—Added VRF option	
Example	`switch (config) # file image upload centos.img.gz scp://username:password@192.168.10.125/var/www/html/<image_name>`	
Related Commands		
Notes		

7.4.16 show docker

	show docker Displays docker running state and VRF context.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.9.2000 3.9.2300—Added new output example

Example	`switch (config) # show docker`
	`Dockers state: started`
	`Docker hub VRF: default`
	`Docker log-level: warn`

Related Commands	
Notes	

7.4.17 show docker containers

	show docker containers <container_name> Displays set parameters on containers already running, and containers planned to run in the future.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.8008 3.8.1000: Updated example 3.9.2000: Updated example, adding host-trust option 3.9.2300: Updated example, adding "user-defined variables" and "log-facility" fields

Example	switch (config) # show docker containers cont_example: image : busybox version : latest status : running start point : data-path-ready cpu limit : 0.2 memory limit: 10m labels : - privileges : network, sdk usb mount : enabled host trust : admin log-facility: kern user-defined variables: name1: value1 name2: value2 another_container: image : busybox version : latest status : - start point : init cpu limit : 0.2 memory limit: 10m labels : my_label privileges : network, sdk usb mount : disabled host trust : admin log-facility: kern user-defined variables: name1: value1 name2: value2 OS_SYSTEM_TYPE : MSN2410 OS_VERSION : 3.9.2300 OS_DOCKERD_VRF_CONTEXT : vrf-default OS_DOCKERD_LINUX_VRF_CONTEXT: vrf_vrf-default switch (config) # show docker containers cont_example cont_example: image : busybox version : latest status : running start point : data-path-ready cpu limit : 0.2 memory limit: 10m labels : - privileges : network, sdk usb mount : enabled host trust : admin log-facility: kern user-defined variables: name1: value1 name2: value2 OS_SYSTEM_TYPE : MSN2410 OS_VERSION : 3.9.2300 OS_DOCKERD_VRF_CONTEXT : vrf-default OS_DOCKERD_LINUX_VRF_CONTEXT: vrf_vrf-default
Related Commands	

Notes	• If a container is already started, the status field displays its current status • If a container is configured to run on the next boot, the start point field displays when it will start • If there is a mismatch between the configuration of a running container and its next-boot configuration, two entries for the container are shown with both of the configurations • For running containers, environment variables that are automatically passed to docker container are revealed (i.e., OS_SYSTEM_TYPE, OS_VERSION, OS_DOCKERD_VRF_CONTEXT, OS_DOCKERD_LINUX_VRF_CONTEXT) • If no user-defined variables were configured, "user-defined variables" field is hidden

7.4.18 show docker images

	show docker images Display docker images.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.3520 3.6.2940: Updated example
Example	

```
switch (config) # show docker images

------------------------------------------------------------
Image              Version      Created            Size
------------------------------------------------------------
ubuntu             latest       Less than a secon  117MB
                                d ago
ubuntu-sdk         v1           41 seconds ago     215MB
```

Related Commands	
Notes	

7.4.19 show docker ps

	show docker ps Display docker containers.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.3520 3.6.2940: Updated example
Example	

```
switch (config) # show docker ps

------------------------------------------------------------------------------
Container           Image:Version          Created               Status
------------------------------------------------------------------------------
my_ubuntu_app       ubuntu:latest          56 seconds ago        Up 50 seconds
```

Related Commands	

Notes	This command is available only after Linux dockers are enabled (“no dockers shutdown”)

7.4.20 show docker labels

	show docker labels Displays docker labels.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.4110

Example	`switch (config) # show docker labels`
	`Storage label : label_name1`
	`configured containers list : cont_name2`
	`active containers list : cont_name1`
	
	`Storage label : label_name2`

Related Commands	
Notes	

7.4.21 show docker login

	show docker login Displays docker login.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.9.1600

Example	`switch (config) # show docker login`
	
	`Servers:`
	`https://index.docker.io/v1/`
	`nvcr.io`

Related Commands	docker login
Notes	

7.4.22 show docker stats

	show docker stats [<name>] Displays Linux docker statistics.	
Syntax Description	name	Docker whose stats to display
Default	N/A	
Configuration Mode	Any command mode	

History	3.6.8008 2.9.2300: Added example
Example	switch (config) # show docker stats --- ------ Container CPU % Memory Memory Memory % Block Block Pids Usage Limit IN OUT --- ------ container1 0.00% 952K 1000M 0.09% 0B 0B 1
Related Commands	
Notes	This command is available only after Linux dockers are enabled (“no dockers shutdown”)

8 Telemetry, Monitoring, and Debuggability

- WHAT JUST HAPPENED
- Logging
- Debugging
- Link Diagnostic Per Port
- Signal Degradation Monitoring
- Event Notifications
- Port Mirroring
- sFlow
- Buffer Histograms Monitoring
- Statistics and Alarms
- Management Information Bases (MIBs)

8.1 WHAT JUST HAPPENED

NVIDIA® WHAT JUST HAPPENED® is based on the extended telemetry capabilities of the NVIDIA Spectrum family switches. This feature, enabled by default, provides the ability to retain the last packets that were dropped from the switch with complete packet headers and the actual drop reason. This enhances the ability to debug network problems, identify affected flows, and decrease time-to-repair.

Retrieving WJH information is done by specifically requesting the last N (up to max 1024 packets per drop reason group) dropped packets & their respective drop reasons. The information is displayed with important Ethernet, IP, and L4 headers. For complete packets, a pcap file is available.

WJH also provides aggregation record for respective drop reasons.

There are four major interfaces enabling the usage of:

- NVIDIA Onyx CLI
- NVIDIA Onyx Web UI
- NEO
- TIG Stack

The following chapters will explain how to use WJH in each of the above modes.

> WJH is only supported through CLI, WebUI, or using NEO, but not in parallel.

8.1.1 Configure What Just Happened (WJH) Using CLI

By Default, NVIDIA® WHAT JUST HAPPENED® is enabled. If it is disabled, use the following command to enable it:

```
switch (config) # what-just-happened <all | acl | forwarding | layer-1 | buffer> enable
```

In Spectrum systems, in order to enable buffer drop monitoring, one interface must be enabled as a recirculation port. For more information see Ethernet Interface Commands section.

To disable WJH via CLI use the “no” form of the command:

```
switch (config) # no what-just-happened <all | acl | forwarding | layer-1 | buffer> enable
```

To display the WJH buffer of dropped packets use the "show what-just-happened" with/without options (detailed in the commands section).

Dropped packet events' display may have a delay of to up to 30 seconds due to a predefined hardware polling interval.

To manually clear WJH buffer use the following command:

```
switch (config) # clear what-just-happened <all | acl | forwarding | layer-1 | buffer>
```

To display the WJH aggregation record, use the "show what-just-happened aggregated" with options (detailed in the commands section).

Note that due to hardware polling timing issues, it may be possible to observe dropped packet events that occurred shortly before the clear command was executed.

To automatically generate a WJH PCAP file as a result of discards, the following configuration is required. The value of <sec> determines how often the system checks whether a pcap should be generated. For example, if you enter a value of 300, up to 5 minutes may elapse between the discarding of packets and the creation of the pcap file.

```
switch (config) # what-just-happened auto-export all enable
switch (config) # logging events what-just-happened-packets enable
switch (config) # logging events what-just-happened-packets interval <sec>
```

To see what pcap files have been generated, issue the following command:

```
switch (config) # show files tcpdump
wjh_auto_export_all_2020_05_18_09_36_12.pcap
```

WJH Wireshark dissector enables Wireshark users to analyze WJH pcap files. It displays the packets' added metadata. You may log into the WebUI and click the "Download Wireshark Plugin" button in the Status → What Just Happened page in order to download the Wireshark plugin file. After downloading the file, place it in the Wireshark application in Windows under %APPDATA%\Wireshark\plugins.

Wireshark dissector was tested on version 2.6.3.

8.1.1.1 WJH Commands

8.1.1.1.1 what-just-happened

	what-just-happened <all | acl | forwarding | layer-1 | buffer> [all-severities | notice | warning | error] enable no what-just-happened <all | acl | forwarding | layer-1 | buffer> [all-severities | notice | warning | error] enable Enables showing dropped packet information. The no form of the command disables showing dropped packet information.	
Syntax Description	all	Drop group containing all packets dropped
	acl	Access-list drops
	forwarding	Drop group containing L2, L3, port and tunnel related drops
	layer-1	Drop group containing layer-1 event
	buffer	Buffer overflow drops
	all-severities	Configure drop with any severity Default: all-severities enabled
	notice	Configure drop with notice severity
	warning	Configure drop with warning severity
	error	Configure drop with error severity
Default	Enabled	
Configuration Mode	config	
History	3.7.1000	
	3.7.1100	Updated Example and Default
	3.8.1000	Updated Syntax and Example
	3.8.2000	Added ACL option
	3.9.0300	Added layer-1 option
	3.9.0500	Added buffer drop option
	3.9.1000	Updated note
	3.9.2000	Updated example and notes, adding support for WJH event suppression by the severity for each drop group.
Example	`switch (config) # what-just-happened forwarding notice enable`	
Related Commands	show what-just-happened status interface ethernet recirculation	

Notes	• In Spectrum systems, in order to enable buffer drop monitoring, one interface must be enabled as a recirculation port. In Spectrum-2 systems, it is sufficient to configure what-just-happened buffer enable. In both cases, the enabling configuration reduces by 1 the number of monitor sessions that can be configured. It will fail if the maximum number of monitor sessions are already configured. • Layer-1 drop group do not support severities • Disabling and enabling the WJH or any drop group will not affect the severity configuration

8.1.1.1.2 what-just-happened auto-export

	what-just-happened auto-export <all | acl | forwarding | buffer> enable no what-just-happened auto-export <all | acl | forwarding | buffer> enable Enables auto-generated pcap files. The no form of the command disables the auto-generation of pcap files.	
Syntax Description	all	Drop group containing all packets dropped
	acl	Access-list drops
	forwarding	Drop group containing L2, L3, port and tunnel related drops
	buffer	Buffer overflow drops
Default	Enabled	
Configuration Mode	configure terminal	
History	3.8.1000	
	3.8.2000	Added ACL option
	3.9.0500	Added buffer drop option
Example	`switch (config) # what-just-happened auto-export forwarding enable`	
Related Commands	what-just-happened enable	
Notes	If auto-export is disabled for acl, forwarding or buffer, dropped packets in those groups do not count towards the threshold for generating a pcap, as defined in the 'logging events what-just-happened-packets' commands.	

8.1.1.1.3 clear what-just-happened

	clear what-just-happened <all | acl | forwarding | layer-1 | buffer> Flushes data from cache DB.	
Syntax Description	all	Drop group containing all packets dropped
	acl	Access-list drops
	forwarding	Drop group containing L2, L3, port and tunnel related drops
	layer-1	Drop group containing layer-1 event
	buffer	Buffer overflow drops
Default	N/A	
Configuration Mode	config	
History	3.7.1000	

	3.8.1000	Updated Syntax and Example
	3.8.2000	Added ACL option
	3.9.0300	Added layer-1 option
	3.9.0500	Added buffer drop option
Example	`switch (config) # clear what-just-happened forwarding`	
Related Commands		
Notes	Clear WJH intends to clear all the events already seen by the user, but will not clear events in the hardware that were not yet read by WJH-lib. As such, it is possible that WJH events observed after using the clear command, actually entered before clearing the command but that were not yet shown to the user.	

8.1.1.1.4 clear what-just-happened pcap-files

	clear what-just-happened pcap-files [all | user | auto-export] Deletes what-just-happened pcap files.	
Syntax Description	all	All PCAP files
	auto-export	PCAP files with wjh_auto_export prefix
	user	PCAP files with wjh_user prefix
Default	all pcap files	
Configuration Mode	config	
History	3.8.2000	
Role	Admin	
Example	`switch (config) # clear what-just-happened pcap-files user`	
Related Commands	file tcpdump delete	
Notes	• All—all pcap files. • User—pcap files with wjh_user prefix. • Auto—exportpcap files with wjh_auto_export prefix.	

8.1.1.1.5 snmp-server notify event what-just-happened

	snmp-server notify event what-just-happened [interval <interval>] [max-traps <max-traps>] no snmp-server notify event what-just-happened [interval <interval>] [max-traps <max-traps>] Enables sending SNMP traps for what-just-happened last events, sets the interval in which traps will be issued, and limits the maximum number of issued traps per interval. The no form of the command disables sending SNMP traps for what-just-happened last events.

Syntax Description	interval	The interval in which traps will be issued. Default: 60 seconds Max: 300 seconds Min: 30 seconds
	max-traps	The maximum number of issued traps per interval. Default: 50 events Max: 100 events Min: 5 events
Default	Disabled	
Configuration Mode	config	
History	3.9.2000	
Example	`switch (config) # snmp-server notify event what-just-happened interval 30 max-traps 100`	
Related Commands	show snmp events what-just-happened	
Notes	• In case SNMP traps for what-just-happened are enabled while using the CLI, a notification will appear informing that SNMP is running in parallel and of the what-just-happened buffer clearing • this command is only relevant for "aggregated" What-Just-Happened events	

8.1.1.1.6 show what-just-happened

	show what-just-happened [all | acl | forwarding | buffer | layer-1 | max-packets <1-1024 per group/1-4096 for all> | export <file-name> | no-metadata] Displays dropped packets information.	
Syntax Description	acl	Access-list drops
	forwarding	Drop group containing L2, L3, port and tunnel related drops
	buffer	Buffer overflow drops
	layer-1	Drop group containing layer 1 event
	max-packets	Limit number of packets to dump: <1-1024> for forwarding/acl/buffer/layer-1, <1-4096> for all Default: 1024 per group, 4096 for all
	export	Create a pcap file
	file name	Optional file-name for pcap file Default: wjh_user_[group]_[date].pcap
	no-metadata	Do not add metadata to the pcap file (applicable only with 'export' attribute set)
Default	N/A	
Configuration Mode	Any command mode	
History	3.7.0000	
	3.7.1100	Updated syntax and example

3.8.1000	Updated syntax, default, and example
3.8.2000	New ACL example
3.8.2100	Update example
3.9.0300	Updated example
3.9.0500	Added layer-1 and buffer drops. PCAP file will not be created by default and updated example
3.9.0900	Updated ACL example
3.9.2000	Updated example of show what-just-happened buffer

Example

```
switch (config) # show what-just-happened

-------------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------------------------------------
------------------------
#     Timestamp                sPort      dPort      VLAN  sMAC               dMAC               EthType
Src IP:Port           Dst IP:Port           IP Proto  Drop Group   Severity  Drop Reason - Recommended
Action
-------------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------------------------------------
------------------------
1     2020/03/31 16:19:51.075  Eth1/3     N/A        12    ba:1b:25:11:22:31  24:8a:07:ca:cd:c8  IPv4
10.10.10.0:6857       10.10.20.1:767        TCP       Forwarding   Warning   Blackhole route - Validate
routing table for this

                     (phonebook)                                          destination IP
...
```

Example (acl)

```
switch (config) # show-what-just-happened acl

-------------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------------------------------------
---------
#     Timestamp                sPort      dPort      VLAN  sMAC               dMAC               EthType
Src IP:Port       Dst IP:Port       IP Proto  Drop Group   Severity  Drop Reason - Recommended
Action
-------------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------------------------------------
---------
1     2020/05/07 12:25:02.600  Eth1/3     N/A        N/A   ba:1b:25:0a:0a:0b  ba:1b:25:0B:0f:01  LPBK
N/A:N/A           N/A:N/A           N/A       Access-list  Notice    user-access-list - Validate ACL
configuration


Rules Info
-------------------------------------------------------------------------------------------------------------------
--------
#     Table Name                                Rule

-------------------------------------------------------------------------------------------------------------------
--------
1     user-access-list                          seq-number 11 deny ba:1b:25:0a:0a:0b mask
ff:ff:ff:ff:ff:ff any


Exception list:
Buffer group is enabled but not operational. Please configure port recirculation.
```

Example (acl export)

```
switch (config) # show what-just-happened acl export

Pcap file path : /vtmp/wjh-pcaps/wjh_user_acl_2020_02_20_11_05_55.pcap

-------------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------
# Timestamp              sPort  dPort VLAN  sMAC              dMAC              EthType   Src IP:Port
Dst IP:Port IP Proto Drop Group Severity Drop Reason - Recommended Action
-------------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------
1 2020/02/20 11:03:17.465 Eth1/3  N/A   N/A  ba:1b:25:0a:0a:0a ba:1b:25:0b:0b:0b  LPBK      N/A:N/A
N/A:N/A      N/A     Access-list  Notice    mac-acl - Validate ACL configuration
```

Example (Layer-1)

```
switch (config) # show what-just-happened layer-1

-------------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------
# Timestamp              sPort  dPort VLAN  sMAC dMAC EthType  Src IP:Port  Dst IP:Port  IP Proto  Drop
Group Severity  Drop Reason - Recommended Action
-------------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------

1 2020/03/16 12:10:58.728 Eth1/15 N/A  N/A   N/A  N/A  N/A      N/A:N/A      N/A:N/A      N/A
Layer-1    Warning   General L1 event - Check layer 1 aggregated

information
```

show what-just-happened all export

```
show (config) # show what-just-happened all export wjh_example

Pcap file path : /vtmp/wjh-pcaps/wjh_ example _all_2020_01_26_10_44_55.pcap

-------------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------------------------
# Timestamp              sPort  dPort VLAN  sMAC              dMAC              EthType  Src IP:Port
Dst IP:Port          IP Proto Drop Group   Severity Drop Reason - Recommended Action
-------------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------------------------
1 2020/01/26 10:44:29.810 Eth1/1  N/A  10    ba:1b:25:11:22:31  24:8a:07:ca:cd:c8  IPv4
10.10.10.0:54401  10.10.20.1:80 (http) TCP                  ACL     Info Openflow Table 1 - Check
Openflow Rule
2 2020/01/26 11:44:29.810 Eth1/2  N/A  20    ee:2b:85:61:22:31  11:2e:ff:ca:cd:d3  IPv4
20.20.20.0:10001  10.10.20.1:80 (http) TCP                  ACL      Info mac-acl - Check ACL Rule

Rules Info
-----------------------------------------------------------------------
# Table Name                          Rule
-----------------------------------------------------------------------
1          Openflow Table 1          ip,ip_dst=10.10.20.1/32,priority=77
2          mac-acl                             seq-number 10 deny ee:2b:85:61:22:31 mask
ff:ff:ff:ff:ff:ff any
```

show what-just-happened buffer

```
switch (config) # show what-just-happened layer-1

-------------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------
# Timestamp               sPort  dPort VLAN  sMAC dMAC EthType  Src IP:Port  Dst IP:Port  IP Proto  Drop
Group Severity  Drop Reason - Recommended Action
-------------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------
1 2020/03/16 12:10:58.728 Eth1/15 N/A  N/A   N/A  N/A  N/A      N/A:N/A      N/A:N/A      N/A
Layer-1    Warning   General L1 event - Check layer 1 aggregated

information
```

show what-just-happened all export

```
show (config) # show what-just-happened all export wjh_example

Pcap file path : /vtmp/wjh-pcaps/wjh_ example _all_2020_01_26_10_44_55.pcap

-------------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------------------------
# Timestamp               sPort  dPort VLAN  sMAC               dMAC               EthType  Src IP:Port
Dst IP:Port          IP Proto Drop Group   Severity Drop Reason - Recommended Action
-------------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------------------------
1 2020/01/26 10:44:29.810 Eth1/1  N/A  10    ba:1b:25:11:22:31  24:8a:07:ca:cd:c8  IPv4
10.10.10.0:54401  10.10.20.1:80 (http) TCP                  ACL     Info Openflow Table 1 - Check
Openflow Rule
2 2020/01/26 11:44:29.810 Eth1/2  N/A  20    ee:2b:85:61:22:31  11:2e:ff:ca:cd:d3  IPv4
20.20.20.0:10001  10.10.20.1:80 (http) TCP                  ACL      Info mac-acl - Check ACL Rule

Rules Info
-------------------------------------------------------------------------
# Table Name                          Rule
-------------------------------------------------------------------------
1          Openflow Table 1          ip,ip_dst=10.10.20.1/32,priority=77
2          mac-acl                             seq-number 10 deny ee:2b:85:61:22:31 mask
ff:ff:ff:ff:ff:ff any
```

show what-just-happened buffer

```
switch (config) # show what-just-happened buffer
-------------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------------------------------------
-----------------
#     Timestamp                sPort      dPort      VLAN  sMAC               dMAC               EthType
Src IP:Port           Dst IP:Port           IP Proto  Drop Group   Severity  Drop Reason - Recommended
Action
-------------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------------------------------------
-----------------
1     2020/10/05 12:23:12.464  Eth1/4     Eth1/2     N/A   98:03:9b:82:bF:7a  b8:59:9f:a6:69:88  IPv4
10.1.2.2:50876        10.1.1.2:2221         (null)    Buffer       Notice    Port TC Congestion Threshold
Crossed - Monitor

                     (rockwell-csp1)                                          network congesti
2     2020/10/05 12:23:12.459  Eth1/4     Eth1/2     N/A   98:03:9b:82:bF:7a  b8:59:9f:a6:69:88  IPv4
   10.1.2.2:50876        10.1.1.2:2221         (null)    Buffer       Notice    Port TC Congestion
Threshold Crossed - Monitor

                    (rockwell-csp1)                                         network congestion
3     2020/10/05 12:23:12.448  Eth1/4     Eth1/2     N/A   98:03:9b:82:bF:7a  b8:59:9f:a6:69:88  IPv4
10.1.2.2:50876        10.1.1.2:2221         (null)    Buffer       Notice    Port TC Congestion Threshold
Crossed - Monitor

                     (rockwell-csp1)                                          network congestion
Buffer Info:

-------------------------------------------------------------------------------------------------------------------
-
  #     TC Id     TC Usage [KB]    Latency [nanoseconds]  TC Watermark [KB]  Latency Watermark
[nanoseconds]

-------------------------------------------------------------------------------------------------------------------
-
  1     1         2896             N/A                    N/A                N/A
  2     1         2960             N/A                    N/A                N/A
  3     1         2920             N/A                    N/A                N/A
```

Related Commands	show what-just-happened status
Notes	• By default, pcap file will not be created, if "export" is not specified. Pcap file names will be “wjh_user_[date].pcap” if no user-defined name is entered, and “[user defined name]_[date].pcap” if provided • In Spectrum systems, in order to see buffer drops, one interface must be configured as a recirculation port • "max-num" and "last-read" are reserved and cannot be used as filenames • For display of ACL drops, lines indexes in "Rules Info" table match the indexes in the main table • To display buffer drops, lines indexes in "Buffer Info" table should match the indexes in the main table

8.1.1.1.7 show what-just-happened aggregated

	show what-just-happened aggregated <acl | forwarding | buffer | layer-1> <max-num | last-read> Displays aggregation record.	
Syntax Description	max-num	Maximum number of aggregated record displayed.
	last-read	Get aggregated record in last show command
	forwarding	Display aggregated record on layer-2/port/layer-3/tunneling related reasons. Max-num is 192.
	acl	Display aggregated record on access list related reasons. Max-num is 64.
	buffer	Max: 64
	layer-1	Display aggregated record on layer-1 related reasons. Max-num is 256.

Default	N/A	
Configuration Mode	Any command mode	
History	3.9.0300	
	3.9.0500	Added support for ACL, forwarding, and buffer drops
	3.9.0900	Updated ACL example
	3.9.1000	In "show-what-just-happened aggregated acl" added support for description of ACL OpenFlow drops
	3.9.2000	Updated example of show what-just-happened aggregated buffer

Example (Layer-1)

```
switch (config) # show what-just-happened aggregated layer-1

Sample Window : 2020/03/19 05:12:54.086 - 2020/03/19 05:47:43.426

-------------------------------------------------------------------------------------------------
 Port State Down Reason - Recommended Action      State Change Symbol Error FCS Error
 ------------------------------------------------------------------------------------------------
 Eth1/4     Down Port admin down - Validate port  1            0            0
            configuration
```

Example (acl)

```
switch (config) # sh what-just-happened aggregated acl

Sample Window : 2020/05/11 10:25:09.953 - 2020/05/11 10:55:11.921
-------------------------------------------------------------------------------------------------------------------
-----------------------------------------------------------
#     sPort     VLAN  sMAC               dMAC               EthType  Src IP:Port    Dst IP:Port    IP
Proto  Count   Severity  Drop Reason - Recommended Action
-------------------------------------------------------------------------------------------------------------------
-----------------------------------------------------------
1     Eth1/2    N/A   ba:1b:25:0a:0a:0a  ba:1b:25:0b:0b:0b  LPBK     N/A:N/A        N/A:N/A         N/
A       2      Notice    mac-acl - Validate ACL configuration


Rules Info
-------------------------------------------------------------------------------------------------------------------
--------
#     Table Name
Rule
-------------------------------------------------------------------------------------------------------------------
--------
1     mac-acl                                 seq-number 10 deny ba:1b:25:0a:0a:0a mask ff:ff:ff:ff:ff:ff
any
```

Example (forwarding)

```
switch (config) # show what-just-happened aggregated forwarding

Sample Window : 2020/03/16 12:10:29.226 - 2020/03/17 02:33:06.337

-------------------------------------------------------------------------------------------------------------------
-----------------------------------------------------------------
# sPort  VLAN sMAC              dMAC              EthType Src IP:Port                    Dst IP:Port
IP Proto Count Severity Drop Reason - Recommended Action
-------------------------------------------------------------------------------------------------------------------
-----------------------------------------------------------------
1 Eth1/2 N/A  24:8a:07:97:32:e2 33:33:00:00:00:16 IPv6    [fe80::268a:7ff:fe97:32e2]:N/A [ff02::16]:N/A
N/A     2      Notice Ingress spanning tree filter- Expected behavior
```

Example (buffer)

```
switch (config) # show what-just-happened aggregated buffer last-read

Sample Window : 2020/10/05 12:36:49.369 - 2020/10/05 12:38:13.605
-------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------
#     sPort     VLAN  sMAC              dMAC              EthType  Src IP:Port           Dst
IP:Port           IP Proto  Count     Severity  Drop Reason - Recommended Action
-------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------
1     Eth1/4    N/A   98:03:9b:82:bF:7a  b8:59:9f:a6:69:88  IPv4     10.1.2.2:50888
10.1.1.2:2221         TCP       6330      Notice    Port TC Congestion Threshold Crossed - Monitor network
                                                                                        (rockwell-
csp1)                                   congestion
Buffer Info:

-------------------------------------------------------------------------------------------------------------
-
  #     TC Id     TC Usage [KB]    Latency [nanoseconds]  TC Watermark [KB]  Latency Watermark
[nanoseconds]

-------------------------------------------------------------------------------------------------------------
-
  1     N/A       N/A              N/A                    371                N/A
```

Related Commands	
Notes	• For display of ACL drops, lines indexes in "Rules Info" table match the indexes in the main table • To display buffer drops, lines indexes in "Buffer Info" table should match the indexes in the main table

8.1.1.1.8 show what-just-happened status

	Show general what-just-happened status Enables and disables what-just-happened drop-groups status.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.8.2000	
	3.9.0600	Added operational status
	3.9.1300	Updated output
	3.9.2000	Updated example, adding "Enabled Severities" and "Operational status" fields
Role	Admin	

Example	switch (config) # show what-just-happened status What-just-happened is enable Severities: N: Notice W: Warning E: Error --- Drop group Admin status Enabled Severities Operational status Auto-export status --- Forwarding Enable W, E Enable Enable Access-list Enable All Enable Enable Buffer Enable N Enable Enable Layer-1 Enable All Enable N/A
Related Commands	
Notes	

8.1.1.1.9 show snmp events what-just-happened

	show snmp events what-just-happened Displays what-just-happened SNMP configuration and status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.9.2000
Example	switch (config) # show snmp events what-just-happened Event : Periodic SNMP traps reporting what-just-happened aggregated data Trap enabled : Yes Polling interval (sec) : 30 Maximum traps per interval: 100 switch (config) # show snmp events what-just-happened Event : Periodic SNMP traps reporting what-just-happened aggregated data Trap enabled : No Polling interval (sec) : N/A Maximum traps per interval: N/A
Related Commands	snmp-server notify event what-just-happened
Notes	

8.1.1.2 Configure WJH Events

8.1.1.2.1 switchmode exceptions sip-equals-dip

	switchmode exceptions sip-equals-dip no switchmode exceptions sip-equals-dip Cancels the discard on this trap. The no form of the command returns the action to discard for this trap.
Syntax Description	N/A
Default	Disabled
Configuration Mode	Config
History	3.9.0900
Example	`switch (config) # switchmode exceptions sip-equals-dip`
Related Commands	show switchmode exceptions
Notes	

8.1.1.2.2 show switchmode exceptions

	show switchmode exceptions Shows the current state of the exceptions
Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.9.0900
Example	`switch (config) # show switchmode exceptions` `Src & Dest IP equal action: Forward`
Related Commands	switchmode exceptions sip-equals-dip
Notes	

8.1.2 Configure WJH Using NEO

For further information of how to install What Just Happened using NEO on an Onyx switch, refer to Installing What Just Happened Using NEO on an Onyx Switch in the NEO User Manual.

8.1.3 WJH Streaming and Integration with Telegraf, InfluxDB and Grafana (TIG) Stack

For further information refer to WJH Streaming and Integration with Telegraf, InfluxDB and Grafana (TIG) Stack in the Telemetry Agent User Manual.

8.2 Logging

8.2.1 Monitor

To print logging events to the terminal, set the modules or events you wish to print to the terminal. For example, run: o-

```
switch (config) # logging monitor events notice
switch (config) # logging monitor sx-sdk warning
```

These commands print system events in severity “notice”, and “sx-sdk” module notifications in severity “warning” to the screen. For example, in case of interface-down event, the following gets printed to the screen:

```
switch (config) #
Wed Jul 10 11:30:42 2013: Interface IB1/17 changed state to DOWN
Wed Jul 10 11:30:43 2013: Interface IB1/18 changed state to DOWN
```

To see a list of the events, refer to “Supported Event Notifications and MIB Mapping”.

8.2.2 Remote Logging

To configure remote syslog to send syslog messages to a remote syslog server:

1. Set remote syslog server.

   ```
   switch (config) # logging <IP address/hostname>
   ```

2. (Optional) Set the destination port of the remote host.

   ```
   switch (config) # logging <IP address/hostname> port <port>
   ```

3. (Optional) Filter log messages according to an input regex.

   ```
   switch (config) # logging <IP address/hostname> filter <"include"/"exclude"> <regex>
   ```

4. Set the minimum severity of the log level to info.

   ```
   switch (config) # logging <IP address/hostname> trap info
   ```

5. Override the log levels on a per-class basis.

   ```
   switch (config) # logging <IP address/hostname> trap override class <class name> priority <level>
   ```

8.2.3 Logging Protocol

A feature that provides the ability to choose the protocol to use for sending syslog messages to a remote host: UDP (default) or TCP. See "logging protocol" command.

8.2.4 Logging Commands

8.2.4.1 logging

	logging [vrf <vrf-name>] <IPv4 address/IPv6 address/hostname> no logging [vrf <vrf-name>] <IPv4 address/IPv6 address/hostname> Sends log messages to the remote host specified by its IP or hostname. The no form of the command stops sending log messages to the remote host specified by its IP or hostname.
Syntax Description	vrf-name—VRF to be affected. If "vrf-name" parameter is not specified, "default" VRF will be used.
Default	N/A
Configuration Mode	config
History	3.1.1000 3.9.2000—Added VRF option
Role	admin
Example	`switch (config) # logging 1.1.1.1` `switch (config) # no logging 1.1.1.1`
Related Commands	
Notes	This command is configurable. If “configuration write” is executed, the remote host will still receive messages after reload. It is possible to have multiple logging hosts in different VRFs.

8.2.4.2 logging port

	logging [vrf <vrf-name>] <syslog IPv4 address/IPv6 address/hostname> port <destination-port> no logging [vrf <vrf-name>] <syslog IPv4 address/IPv6 address/hostname> port Configures remote server destination port for log messages. The no form of the command resets the remote log port to its default value.	
Syntax Description	destination-port	Range: 1-65535
	Hostname	Max 64 characters
	vrf-name-VRF to be affected. If "vrf-name" parameter is not specified, "default" VRF will be used.	
Default	514 (UDP)	
Configuration Mode	config	
History	3.6.2002 3.8.1000—Updated command syntax 3.9.2000—Added VRF option	

Example	`switch (config) # logging 10.0.0.1 port 105`
Related Commands	logging <syslog IPv4 address/IPv6 address/hostname> trap
Notes	It is possible to have multiple logging hosts in different VRFs.

8.2.4.3 logging trap

	logging [vrf <vrf-name>] <syslog IPv4 address/IPv6 address/hostname> [trap {<log-level> | override class <class> priority <log-level>}] no logging [vrf <vrf-name>] <syslog IPv4 address/IPv6 address/hostname> [trap {<log-level> | override class <class> priority <log-level>}] Enables (by setting the syslog IPv4 address/IPv6 address/hostname) sending logging messages, with ability to filter the logging messages according to their classes. The no form of the command stops sending messages to the remote syslog server.	
Syntax Description	syslog IPv4 address/IPv6 address/ hostname	syslog IPv4 address/IPv6 address/hostname of the remote syslog server Hostname is limited to 64 characters
	log-level	• none—disables the logging locally and remotely • 0 - emerg—system is unusable (emergency) • 1 - alert—alert notification, action must be taken immediately • 2 - crit—critical condition • 3 - err—error condition • 4 - warning—warning condition • 5 - notice—normal, but significant condition • 6 - info—informational condition • 7 - debug—debug level messages
	class	Sets or removes a per-class override on the logging level. All classes which do not have an override set will use the global logging level set with “logging local <log level>”. Classes that do have an override will do as the override specifies. If “none” is specified for the log level, the software will not log anything from this class. Classes available: • iss-modules—protocol stack • mgmt-back—system management back-end • mgmt-core—system management core • mgmt-front—system management front-end • mlx-daemons—management daemons • sx-sdk—switch SDK
	vrf-name	VRF to be affected. If "vrf-name" parameter is not specified, "default" VRF will be used.
Default	Remote logging is disabled	
Configuration Mode	config	
History	3.6.2002 3.8.1000—Updated command syntax 3.9.2000—Added VRF option	
Example	`switch (config) # logging local info`	
Related Commands	show logging logging local override logging <syslog IPv4 address/IPv6 address/hostname> port	
Notes	It is possible to have multiple logging hosts in different VRFs.	

8.2.4.4 logging debug-files

	logging debug-files {delete {current | oldest} | rotation {criteria | force | max-num} | update {<number> | current} | upload <log-file> <upload URL>} no logging debug-files rotation criteria Configures settings for debug log files. The "logging debug-files rotation criteria" command removes the debug rotation criteria configuration.	
Syntax Description	delete {current | oldest}	Deletes certain debug-log files. • current—deletes the current active debug-log file • oldest—deletes some of the oldest debug-log files

	rotation {criteria {frequency {daily | weekly | monthly} | size <size> | size-pct <percentage>} | force | max-num}	Configures automatic rotation of debug-logging files.
		• criteria—sets how the system decides when to rotate debug files
		  • frequency—rotate log files on a fixed time-based schedule
		  • size—rotate log files when they pass a size threshold in megabytes
		  • size-pct—rotate logs when they surpass a specified percentage of disk
		• forces—forces an immediate rotation of the log files
		• max-num—specifies the maximum number of old log files to keep

	update {<number> | current}	Uploads a local debug-log file to a remote host. • current—uploads log file “messages” to a remote host • number—uploads compressed log file “debug.<number>.gz” to a remote host. Range is 1-10.
	upload	Uploads debug log file to a remote host
	log-file	Possible values: 1-7, or current
	upload URL	Supported formats: HTTP, HTTPS, FTP, TFTP, SCP and SFTP (e.g.: scp://username[:password]@hostname/path/filename)
Default	N/A	
Configuration Mode	config	
History	3.3.4150 3.9.0900: Added "no logging debug-files rotation criteria" command	
Example	`switch (config) # logging debug-files delete current`	
Related Commands		
Notes		

8.2.4.5 logging events enable

	logging events {cpu-rate-limiters | interfaces | protocols | what-just-happened-packets} enable no logging events {cpu-rate-limiters | interfaces | protocols | what-just-happened-packets| what-just-happened-packets} enable Activate event tracking for a certain group. The no form of the command deactivates event tracking for a certain group.

Syntax Description	cpu-rate-limiters | interfaces | protocols | what-just-happened-packets	Logical groups with specified set of counters
Default	N/A	
Configuration Mode	config	
History	3.6.6000 3.9.0900: Added note and what-just-happened-packets option	
Example	`switch (config) # logging events interfaces enable`	
Related Commands		

Notes	`Increase in the enabled events groups will generate a log message of the form:`
	`Jan 8 14:15:24 switch statsd[4404]: [statsd.NOTICE]: (StatsLog) Interface Eth1/9: 398 0598 packets dropped due to Rx invalid tag discards packets`
	`Jan 8 14:15:24 switch statsd[4404]: [statsd.NOTICE]: (StatsLog) Interface Eth1/9: 398 0599 packets dropped due to Rx discard packets by vlan filter`
	`Jan 8 14:42:44 switch statsd[4404]: [statsd.NOTICE]: (StatsLog) cpu-rate-limiter DISCARD_LAYERS_2_3: 7767087 packets dropped by CPU rate-limiter`	

8.2.4.6 logging events error-threshold

	logging events {cpu-rate-limiters | interfaces | protocols | what-just-happened-packets} error-threshold <events> no logging events {cpu-rate-limiters | interfaces | protocols | what-just-happened-packets} error-threshold <events> Configures number of events after which the system begins to generate events to the log file. The no form of the command resets this parameter to its default value.	
Syntax Description	cpu-rate-limiters	Sets threshold for CPU rate limiter related events Default: 1 event
	interfaces	Sets threshold for interface related events Default: 10 events
	protocols	Sets threshold for protocol related events Default: 2 events
	what-just-happened-packets	Sets threshold for dropped packets Default: 1000 packets
	events	Number of events after which the system begins to generate events to the log file. Range: 0-4294967295.
Default	N/A	
Configuration Mode	config	
History	3.6.6000	
	3.9.0900	Added what-just-happened-packets options and note
Example	`switch (config) # logging events interfaces error-threshold 45`	

Related Commands	
Notes	The command configures number of events after which the system begins to generate events to the log file, if that number of events occurs within the interval defined by the logging events interval command. In the case of what-just-happened-packets, number of events refers to the number of dropped packets due to reasons for which auto-export pcap generation is enabled.

8.2.4.7 logging events interval

	logging events {cpu-rate-limiters | interfaces | protocols| what-just-happened-packets} interval <seconds> no logging events {cpu-rate-limiters | interfaces | protocols | what-just-happened-packets} interval <seconds> Configures interval in seconds between each sampling of counters in event type. The no form of the command resets this parameter to its default value.	
Syntax Description	cpu-rate-limiters | interfaces | protocols | what-just-happened-packets	Logical groups with specified set of counters Default:cpu-rate-limiters—10 seconds interfaces—5 minutes protocols—1 minute what-just-happened-packets - 10 minutes
	seconds	Time between sampling. Range is different for each event type: • cpu-rate-limiters—5-3600 • interfaces—10-3600 • protocols—10-3600 • what-just-happened-packets—10-3600
Default	N/A	
Configuration Mode	config	
History	3.6.6000 3.9.0900: Added what-just-happened-packets option	
Example	`switch (config) # logging events interfaces interval 120`	
Related Commands		
Notes	In the case of what-just-happened-packets, a pcap file will be generated if the threshold number of dropped packets is exceeded during this interval	

8.2.4.8 logging events rate-limit

	logging events [cpu-rate-limiters | interfaces | protocols | what-just-happened-packets] rate-limit {short | medium | long} [count | window] no logging events [cpu-rate-limiters | interfaces | protocols | what-just-happened-packets] rate-limit [short | medium | long] [count <number> | window <seconds>] Configures the number of allowed events per time window and that window's duration. The no form of the command resets these parameters to their default values.

Syntax Description	cpu-rate-limiters | interfaces | protocols | what-just-happened-packets	Logical groups with specified set of counters		
	rate-limit	Three configurable periods: short, medium, and long		
	count	Number of allowed events per time window		
	window	Window of time in seconds for the rate limit period		
Default	For “interfaces” Short window: event count—5 window duration—1 hour Medium window: event count—50 window duration—1 day Long window: event count—350 window duration—7 days	For “protocols” Short window: event count—10 window duration—1 hour Medium window: event count—100 window duration—1 day Long window: event count—600 window duration—7 days	For “cpu-rate-limiters” Short window: event count—10 window duration—1 hour Medium window: event count—200 window duration—1 day Long window: event count—1200 window duration—7 days	For “what-just-happened-packets” Short window: event count—3 window duration—1 hour Medium window: event count—15 window duration—1 day Long window: event count—50 window duration—7 days
Configuration Mode	config			
History	3.6.6000			
	3.9.0900	Added what-just-happened-packets option		
Example	`switch (config) # logging events interfaces interval 120`			
Related Commands				
Notes	• The goal of this command is to restrict the number of events in the log, or, in the case of what-just-happened-packets, the number of pcap files generated. To achieve this end, it is possible to specify the allowed number (parameter “count”) of messages per period of time (parameter “window”) • In the case of what-just-happened-packets, the configured logging events rate-limit configures is the maximum number of pcap files that may be generated in each time window			

8.2.4.9 logging fields

	logging fields seconds {enable | fractional-digits <f-digit> | whole-digits <w-digit>} no logging fields seconds {enable | fractional-digits <f-digit> | whole-digits <w-digit>} Specifies whether to include an additional field in each log message that shows the number of seconds since the Epoch or not. The no form of the command disallows including an additional field in each log message that shows the number of seconds since the Epoch.	
Syntax Description	enable	Specifies whether to include an additional field in each log message that shows the number of seconds since the Epoch or not.
	f-digit	The fractional-digits parameter controls the number of digits to the right of the decimal point. Truncation is done from the right. Possible values are: 1, 2, 3, or 6.

	w-digit	The whole-digits parameter controls the number of digits to the left of the decimal point. Truncation is done from the left. Except for the year, all of these digits are redundant with syslog's own date and time. Possible values: 1, 6, or all.
Default	Disabled	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # logging fields seconds enable` `switch (config) # logging fields seconds whole-digits 1`	
Related Commands	show logging	
Notes	This is independent of the standard syslog date and time at the beginning of each message in the format of “July 15 18:00:00”. Aside from indicating the year at full precision, its main purpose is to provide subsecond precision.	

8.2.4.10 logging files delete

	logging files delete {current | oldest [<number of files>]} Deletes the current or oldest log files.	
Syntax Description	current	Deletes current log file
	oldest	Deletes oldest log file
	number of files	Sets the number of files to be deleted
Default	CLI commands and audit message are set to notice logging level	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # logging files delete current`	
Related Commands	show logging show log files	
Notes		

8.2.4.11 logging files rotation

	logging files rotation {criteria {frequency <freq> | size <size-mb>| size-pct <size-percentage>} | force | max-number <number-of-files>} no logging files rotation criteria Sets the rotation criteria of the logging files. The no form of the command removes the rotation criteria configuration.	
Syntax Description	freq	Sets rotation criteria according to time. Possible options are: • Daily • Weekly • Monthly
	size-mb	Sets rotation criteria according to size in megabytes Range: 1-9999 Default: 20MB

	size-percentage	Sets rotation criteria according to size in percentage of the partition where the logging files are kept in. The percentage given is truncated to three decimal points (thousandths of a percent).
	force	Forces an immediate rotation of the log files. This does not affect the schedule of auto-rotation if it was done based on time: the next automatic rotation will still occur at the same time for which it was previously scheduled. Naturally, if the auto-rotation was based on size, this will delay it somewhat as it reduces the size of the active log file to zero.
	number-of-files	The number of log files will be kept. If the number of log files ever exceeds this number (either at rotation time, or when this setting is lowered), the system will delete as many files as necessary to bring it down to this number, starting with the oldest.
Default	10 files are kept by default with rotation criteria of 5% of the log partition size	
Configuration Mode	config	
History	3.1.0000 3.9.0900: • Added the command "no logging files rotation criteria" • Changed default value size from 19.07 MB to 20 MB	
Example	`switch (config) # logging files rotation criteria size-pct 6`	
Related Commands	show logging show log files	
Notes		

8.2.4.12 logging files upload

	logging files upload {current | <file-number>} <url> Uploads a log file to a remote host.	
Syntax Description	current	The current log file. The current log file will have the name “messages” if you do not specify a new name for it in the upload URL.
	file-number	An archived log file. The archived log file will have the name “messages<n>.gz” (while “n” is the file number) if you do not specify a new name for it in the upload URL. The file will be compressed with gzip.
	url	Uploads URL path. Supported formats: FTP, TFTP, SCP, and SFTP. For example: scp://username[:password]@hostname/path/filename.
Default	10 files are kept by default with rotation criteria of 5% of the log partition size	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # logging files upload 1 scp://admin@scpserver`	
Related Commands	show logging show log files	
Notes		

8.2.4.13 logging filter include

	logging <IP address\hostname> filter include <regex> Sends only log messages that match the input regex to a remote host specified by its IP or hostname.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.8.2000
Role	admin
Example	`switch (config) # logging 1.1.1.1 filter include ERROR`
Related Commands	loggin no logging
Notes	This command is configurable. If “configuration write” is executed, the remote host will still receive filtered messages after reload.

8.2.4.14 logging filter exclude

	logging <IP address\hostname> filter exclude <regex> Sends only log messages that do not match the input regex to a remote host specified by its IP or hostname.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.8.2000
Role	admin
Example	`switch (config) # logging 1.1.1.1 filter exclude ERROR`
Related Commands	loggin no logging
Notes	This command is configurable. If “configuration write” is executed, the remote host will still receive filtered messages after reload.

8.2.4.15 no logging filter

	no logging <IP address\hostname> filter Sends unfiltered log messages to the configured remote host.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.8.2000
Role	admin

Example	`switch (config) # no logging 1.1.1.1 filter`
Related Commands	loggin no logging
Notes	This command is configurable. If “configuration write” is executed, the remote host will still receive filtered messages after reload.

8.2.4.16 logging format

	logging format {standard | welf [fw-name <hostname>]} no logging format {standard | welf [fw-name <hostname>]} Sets the format of the logging messages. The no form of the command resets the format to its default.	
Syntax Description	standard	Standard format
	welf	WebTrends Enhanced Log file (WELF) format
	hostname	Specifies the firewall hostname that should be associated with each message logged in WELF format. If no firewall name is set, the hostname is used by default. Hostname is limited to 64 characters.
Default	standard	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # logging format standard`	
Related Commands	show logging	
Notes		

8.2.4.17 logging level

	logging level {cli commands <log-level> | audit mgmt <log-level>} Sets the severity level at which CLI commands or the management audit message that the user executes are logged. This includes auditing of both configuration changes and actions.	
Syntax Description	cli commands	Sets the severity level at which CLI commands which the user executes are logged
	audit mgmt	Sets the severity level at which all network management audit messages are logged
	log-level	• none—disables the logging locally and remotely • 0 - emerg—system is unusable (emergency) • 1 - alert—alert notification, action must be taken immediately • 2 - crit—critical condition • 3 - err—error condition • 4 - warning—warning condition • 5 - notice—normal, but significant condition • 6 - info—informational condition • 7 - debug—debug level messages
Default	CLI commands and audit message are set to notice logging level	
Configuration Mode	config	

History	3.1.0000
Example	`switch (config) # logging level cli commands info`
Related Commands	show logging
Notes	

8.2.4.18 logging local override

	logging local override [class <class> priority <log-level>] no logging local override [class <class> priority <log-level>] Enables class-specific overrides to the local log level. The no form of the command disables all class-specific overrides to the local log level without deleting them from the configuration, but disables them so that the logging level for all classes is determined solely by the global setting.	
Syntax Description	override	Enables class-specific overrides to the local log level.
	class	Sets or removes a per-class override on the logging level. All classes which do not have an override set will use the global logging level set with “logging local <log level>”. Classes that do have an override will do as the override specifies. If “none” is specified for the log level, the software will not log anything from this class. Classes available: • debug-module—debug module functionality • protocol-stack—protocol stack modules functionality • mgmt-back—system management back-end components • mgmt-core—system management core • mgmt-front—system management front-end components • mlx-daemons—management daemons • sx-sdk—switch SDK
	log-level	• none—disables the logging locally and remotely • 0 - emerg—system is unusable (emergency) • 1 - alert—alert notification, action must be taken immediately • 2 - crit—critical condition • 3 - err—error condition • 4 - warning—warning condition • 5 - notice—normal, but significant condition • 6 - info—informational condition • 7 - debug—debug level messages
Default	Override is disabled	
Configuration Mode	config	
History	3.1.0000 3.3.4150: Added debug-module class and changed iss-modules to protocol-stack	
Example	`switch (config) # logging local override class mgmt-front priority warning`	
Related Commands	show logging logging local	
Notes		

8.2.4.19 logging monitor

	logging monitor <facility> <priority-level> no logging monitor <facility> <priority-level> Sets monitor log facility and level to print to the terminal. The no form of the command disables printing logs of facilities to the terminal.	
Syntax Description	facility	• mgmt-front • mgmt-back • mgmt-core • events • sx-sdk • mlnx-daemons • iss-modules
	priority-level	• none • emerg • alert • crit • err • warming • notice • info • debug
Default	no logging monitor	
Configuration Mode	config	
History	3.3.4000	
Example	`switch (config) # logging monitor events notice`	
Related Commands		
Notes		

8.2.4.20 logging protocol

	logging <IP address\hostname> protocol [tcp|udp] no logging <IP address\hostname> protocol Sends log messages to specified host with the chosen protocol (TCP or UDP). The no form of the command sets the protocol for sending log messages to a remote host to the default (UDP).	
Syntax Description	tcp	Sets protocol to TCP
	udp	Sets protocol to UDP
Default	UDP	
Configuration Mode	Configure terminal	
History	3.8.2100	
Role	Admin	
Example	`switch (config) # logging 1.1.1.1 protocol tcp` `switch (config) # no logging 1.1.1.1 protocol`	
Related Commands		

Notes	This command is configurable, so if “configuration write” is executed then after reboot the remote host will still receive messages with the configured protocol.

8.2.4.21 logging receive

	logging receive no logging receive Enables receiving logging messages from a remote host. The no form of the command disables the option of receiving logging messages from a remote host.
Syntax Description	N/A
Default	Receiving logging is disabled
Configuration Mode	config
History	3.1.0000
Example	`switch (config) # logging receive`
Related Commands	show logging logging local logging local override
Notes	• This does not log to the console TTY port • In-band management should be enabled in order to open a channel from the host to the CPU • If enabled, only log messages matching or exceeding the minimum severity specified with the “logging local” command will be logged, regardless of what is sent from the remote host

8.2.4.22 logging mac masking

	logging mac masking no logging mac masking Enables MAC address masking in logs. The no form of the command disables MAC address masking.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config
History	3.9.0900
Example	`switch (config) # logging mac masking`
Related Commands	show logging
Notes	If enabled, the first 2 bytes of MAC address output log will be masked. For example, 00:12:34:56:78:9a will be displayed as **:**:34:56:78:9a.

8.2.4.23 show log

	show log [continuous | files [<file-number>]] [[not] matching <reg-exp>] Displays the log file with optional filter criteria.

Syntax Description	continues	Displays the last few lines of the current log file and then continues to display new lines as they come in until the user hits Ctrl+C, similar to LINUX “tail” utility
	files	Displays the list of log files
	<file-number>	Displays an archived log file, where the number may range from 1 up to the number of archived log files available
	[not] matching <reg-exp>	The file is piped through a LINUX “grep” utility to only include lines either matching, or not matching, the provided regular expression
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000 3.3.4402: Updated example and added note	

Example

```
switch (config) # show log matching "Executing|Action"
Jul 31 16:11:23 M2100-aj cli[26502]: [cli.NOTICE]: user : Executing command: enable
Jul 31 16:11:24 M2100-aj cli[26507]: [cli.NOTICE]: user : Executing command: enable
Jul 31 16:11:29 M2100-aj cli[26514]: [cli.NOTICE]: user : Executing command: enable
Jul 31 16:11:29 M2100-aj cli[26514]: [cli.NOTICE]: user : Executing command: show license
Jul 31 16:11:41 M2100-aj cli[26548]: [cli.NOTICE]: user : Executing command: enable
Jul 31 16:11:42 M2100-aj cli[26553]: [cli.NOTICE]: user : Executing command: enable
Jul 31 16:11:42 M2100-aj cli[26553]: [cli.NOTICE]: user : Executing command: conf termina
```

Related Commands	logging fields logging files rotation logging level logging local logging receive show logging
Notes	• When using a regular expression containing | (OR), the expression should be surrounded by quotes (“<expression>”), otherwise it is parsed as filter (PIPE) command • The command’s output has many of the options as the Linux “less” command. These options allow navigating the log file and perform searches. To see help for different option press “h” after running the “show log” command.

8.2.4.24 show logging

	show logging Displays the logging configurations.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.8.2000: Updated example 3.9.0900: Updated example
Role	Admin

Example	switch (config) # show logging Local logging level : notice Override for class debug-module : notice Default remote logging level : notice Allow receiving of messages from remote hosts: no Number of archived log files to keep : 10 Log rotation size threshold : 19.07 megabytes Log rotation (debug) size threshold : 19.07 megabytes Log format : standard Subsecond timestamp field : disabled MAC address masking : enabled Levels at which messages are logged: CLI commands : notice Audit messages: notice Remote syslog servers: 1.1.1.1: log level : notice Remote port : 514 Filter [include] regex: err 1.2.2.3: log level : notice Remote port: 33
Related Commands	logging fields logging files rotation logging level logging local logging receive logging <syslog IPv4 address/IPv6 address/hostname>
Notes	

8.2.4.25 show logging events

	show logging events [cpu-rate-limiters | interfaces | protocols | what-just-happened-packets] Displays configuration per selected event group or all.	
Syntax Description	cpu-rate-limiters | interfaces | protocols | what-just-happened-packets	Logical groups with specified set of counters
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6000	
	3.9.0900	Added what-just-happened-packets option

Example	switch (config) # show logging events
	
	cpu-rate-limiters:
	  Admin mode : yes
	  Interval : 10 seconds
	  Error threshold: 1
	
	  Rate-limit short window:
	  Event count : 10
	  Window duration: 1 hour
	
	  Rate-limit medium window:
	  Event count : 200
	  Window duration: 1 day
	
	  Rate-limit long window:
	  Event count : 1200
	  Window duration: 7 days
	
	interfaces:
	  Admin mode : no
	  Interval : 5 minutes
	  Error threshold: 10
	
	  Rate-limit short window:
	  Event count : 5
	  Window duration: 1 hour
	
	  Rate-limit medium window:
	  Event count : 50
	  Window duration: 1 day
	
	  Rate-limit long window:
	  Event count : 350
	  Window duration: 7 days
	
	protocols:
	  Admin mode : no
	  Interval : 1 minute
	  Error threshold: 2
	
	  Rate-limit short window:
	  Event count : 10
	  Window duration: 1 hour
	
	  Rate-limit medium window:
	  Event count : 100
	  Window duration: 1 day
	
	  Rate-limit long window:
	  Event count : 600
	  Window duration: 7 days
	
	what-just-happened-packets:
	  Admin mode : no
	  Interval : 1 minute
	  Error threshold: 2
	
	  Rate-limit short window:
	  Event count : 10
	  Window duration: 1 hour
	
	  Rate-limit medium window:
	  Event count : 100
	  Window duration: 1 day
	
	  Rate-limit long window:
	  Event count : 600
	  Window duration: 7 days

Related Commands	
Notes	

8.2.4.26 show logging events source-counters

	show logging events [cpu-rate-limiters |interfaces | protocols] source-counters Displays set of counters for sampling.	
Syntax Description	cpu-rate-limiters | interfaces | protocols	Logical groups with specified set of counters
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6000	

Example	`switch (config) # show logging events interfaces source-counters`
	
	`interfaces:`
	`  Counters: Rx discard packets, Rx error packets, Rx fcs errors, Rx undersize packets, Rx oversize packets, Rx unknown control opcode, Rx symbol errors, Rx discard packets by Storm Control, Tx discard packets, Tx error packets, Tx hoq discard packets`	

Related Commands	logging event enable logging event error-threshold logging event interval logging event rate-limit	
Notes		

8.2.4.27 show logging port

	show logging port Displays the port logging configurations.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.8.1000: Updated example

Example	`switch (config) # show logging`
	`Local logging level: notice`
	`  Override for class debug-module: notice`
	`Default remote logging level: notice`
	`Remote syslog receiver: 1.2.3.4 (log level: notice)`
	`Remote port: 514`

Related Commands	logging port
Notes	

8.3 Debugging

To use the debugging logs feature:

1. Enable debugging. Run:

```
switch (config) # debug ethernet all
```

2. Display the debug level set. Run:

```
switch (config) # show debug ethernet
```

3. Display the logs. Run:

```
switch (config) # show log debug {match | continue}
```

8.3.1 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community post:

- HowTo Generate and Upload Debug Dump on NVIDIA Switches
- HowTo Troubleshoot NVIDIA Ethernet Switches via Port Counters

8.3.2 Debugging Commands

8.3.2.1 debug ethernet all

	debug ethernet all no debug ethernet all Enables debug traces for Ethernet modules. The no form of the command disables the debug traces for all Ethernet modules.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.3.4150
Example	`switch (config) # debug ethernet all`
Related Commands	show debug ethernet
Notes	

8.3.2.2 debug ethernet dcbx

	debug ethernet dcbx {all | management | fail-all | control-panel | tlv} Configures the trace level for DCBX. The no form of the command disables the configured DCBX debug traces.	
Syntax Description	all	Enables all traces
	management	Management messages
	fail-all	All failure traces
	control-panel	Control plane traces
	tlv	TLV related trace configuration

Default	N/A
Configuration Mode	config
History	3.3.4150
Example	`switch (config) # debug ethernet dcbx all`
Related Commands	show debug ethernet
Notes	

8.3.2.3 debug ethernet ip igmp-snooping

	debug ethernet ip igmp-snooping {all | forward-db-messages | group-info | init-shut | packet-dump | query | source-info | system-resources-management | timer | vlan-info | filter | max-groups} no debug ethernet ip igmp-snooping {all | forward-db-messages | group-info | init-shut | packet-dump | query | source-info | system-resources-management | timer | vlan-info | filter | max-groups} Configures the trace level for IGMP snooping. The no form of the command disables tracking a specified level.	
Syntax Description	all	Enable track traces
	forward-db-messages	Forwarding database messages
	group-info	Group information messages
	init-shut	Init and shutdown messages
	packet-dump	Packet dump messages
	query	Query related messages
	source-info	Source information messages
	system-resources-management	System resources management messages
	timer	Timer messages
	vlan-info	VLAN information messages
	filter	Filter profile messages
	max-groups	Filter max-groups messages
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
	3.9.2100	Added support for IGMP snooping filtering option (filter and max-groups options)
Example	`switch (config) # debug ethernet ip igmp-snooping all`	
Related Commands	show debug ethernet	
Notes		

8.3.2.4 debug ethernet ip interface

	debug ethernet ip interface {all | arp-packet-dump | buffer | enet-packet-dump | error | fail-all | filter | trace-error | trace-event} no debug ethernet ip interface {all | arp-packet-dump | buffer | enet-packet-dump | error | fail-all | filter | trace-error | trace-event} Configures the trace level for interface. The no form of the command disables tracking a specified level.	
Syntax Description	all	Enable track traces
	arp-packet-dump	ARP packet dump trace
	buffer	Buffer trace
	enet-packet-dump	ENET packet dump trace
	error	Trace error messages
	fail-all	All failures including Packet Validation Trace
	filter	Lower layer traces
	trace-error	Trace error messages
	trace-event	Trace event messages
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
Example	`switch (config) # debug ethernet ip interface all`	
Related Commands	show debug ethernet	
Notes		

8.3.2.5 debug ethernet lacp

	debug ethernet lacp {all | all-resource | data-path | fail-all | init-shut | management | memory | packet} no debug ethernet lacp {all | all-resources | data-path | fail-all | init-shut | management | memory | packet} Configures the trace level for LACP. The no form of the command disables the configured LACP debug traces.	
Syntax Description	all	Enables all traces
	all-resource	BPDU related messages
	data-path	Init and shutdown traces
	fail-all	Management messages
	init-shut	Memory related messages
	management memory	IP packet dump trace
	memory	All failure traces
	packet	OS resource trace

Default	N/A
Configuration Mode	config
History	3.3.4150
Example	`switch (config) # debug ethernet lacp all`
Related Commands	show debug ethernet
Notes	

8.3.2.6 debug ethernet lldp

	debug ethernet lldp {all | control-panel | critical-event | data-path | fail-all | init-shut | management | memory | neigh-add | neigh-age-out | neigh-del | neigh-drop | neigh-updt | tlv} no debug ethernet lldp {all | control-panel | critical-event | data-path | fail-all | init-shut | management | memory | neigh-add | neigh-age-out | neigh-del | neigh-drop | neigh-updt | tlv} Configures the trace level for LLDP. The no form of the command disables the configured LLDP debug traces.	
Syntax Description	all	Enables all traces
	control-panel	Control plane traces
	critical-event	Critical traces
	data-path	IP packet dump trace
	fail-all	All failure traces
	init-shut	Init and shutdown traces
	management	Management messages
	memory	Memory related messages
	neigh-add	Neighbor add traces
	neigh-age-out	Neighbor ageout traces
	neigh-del	Neighbor delete traces
	neigh-drop	Neighbor drop traces
	neigh-updt	Neighbor update traces
	tlv	TLV related trace configuration
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
Example	`switch (config) # debug ethernet lldp all`	
Related Commands	show debug ethernet	
Notes		

8.3.2.7 debug ethernet port

	debug ethernet port all Configures the trace level for port. The no form of the command disables the configured port debug traces.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.3.4150
Example	`switch (config) # debug ethernet port all`
Related Commands	show debug ethernet
Notes	

8.3.2.8 debug ethernet qos

	debug ethernet qos {all | all-resource | control-panel | fail-all | filters | init-shut | management | memory | packet} no debug ethernet qos {all | all-resource | control-panel | fail-all | filters | init-shut | management | memory | packet} Configures the trace level for QoS. The no form of the command disables the configured QoS debug traces.	
Syntax Description	all	Enables all traces
	all-resource	OS resource traces
	control-panel	Control plane traces
	fail-all	All failure traces
	filters	Lower layer traces
	init-shut	Init and shutdown traces
	management	Management messages
	memory	Memory related messages
	packet	BPDU related messages
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
Example	`switch (config) # debug ethernet port all`	
Related Commands	show debug ethernet	
Notes		

8.3.2.9 debug ethernet spanning-tree

	debug ethernet spanning-tree {all | error | event | filters | init-shut | management | memory | packet | port-info-state-machine | port-receive-state-machine | port-role-selection-state-machine | port-transit-state-machine | port-transmit-state-machine | protocol-migration-state-machine | timers} no debug ethernet spanning-tree {all | error | event | filters | init-shut | management | memory | packet | port-info-state-machine | port-receive-state-machine | port-role-selection-state-machine | port-transit-state-machine | port-transmit-state-machine | protocol-migration-state-machine | timers} Configures the trace level for spanning-tree. The no form of the command disables the configured spanning-tree debug traces.	
Syntax Description	all	Enables all traces
	error	Error messages trace
	event	Events related messages
	filters	Lower later traces
	init-shut	Init and shutdown traces
	management	Management messages
	memory	Memory related messages
	packet	BPDU related messages
	port-info-state-machine	Port information messages
	port-receive-state-machine	Port received messages
	port-role-selection-state-machine	Port role selection messages
	port-transit-state-machine	Port transition messages
	port-transmit-state-machine	Port transmission messages
	protocol-migration-state-machine	Protocol migration messages
	timers	Timer modules message
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
Example	`switch (config) # debug ethernet spanning-tree all`	
Related Commands	show debug ethernet	
Notes		

8.3.2.10 debug ethernet vlan

	debug ethernet vlan {all | fwd | priority | filters} no debug ethernet vlan {all | fwd | priority | filters} Configures the trace level for VLAN. The no form of the command disables the configured VLAN debug traces.	
Syntax Description	all	Enables all traces
	fwd	Forward
	priority	Priority
	filters	Lower layer traces
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
Example	`switch (config) # debug ethernet vlan all`	
Related Commands	show debug ethernet	
Notes		

8.3.2.11 show debug ethernet

	show debug ethernet {dcbx | ip {arp | dhcp-relay | igmp-snooping | interface | ospf} | lacp | lldp | port | qos | spanning-tree | vlan} Displays debug level configuration on a specific switch.	
Syntax Description	dcbx	Displays the trace level for spanning tree
	ip	Displays debug trace level for ethernet routing module: • arp • dhcp-relay • igmp-snooping • interface • ospf
	lacp	Displays the trace level for LACP
	lldp	Displays the trace level for LLDP
	port	Displays the trace level for port
	qos	Displays the trace level for QoS
	spanning-tree	Displays the trace level for spanning tree
	vlan	Displays the trace level for VLAN
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4150	
	3.6.6000	Updated Example

Example	switch (config) # show debug ethernet dcbx
	dcbx protocol:
	  management : ON
	  fail-all : ON
	  control-panel: ON
	  tlv : ON

Related Commands	debug ethernet all debug ethernet dcbx debug ethernet ip igmp-snooping debug ethernet ip interface debug ethernet lacp debug ethernet lldp debug ethernet port debug ethernet qos debug ethernet spanning-tree debug ethernet vlan
Notes	

8.3.2.12 show log debug

	show log debug [continuous | files | matching | not] Displays current event debug-log file in a scrollable pager.	
Syntax Description	continuous	Displays new event log messages as they arrive
	files	Displays archived debug log files
	matching	Displays event debug logs that match a given regular expression
	not	Displays event debug logs that do not meet certain criteria
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4150	
Example		

```
switch (config) # show log debug

Jun 15 16:20:47 switch-627e4c last message repeated 7 times
Jun 15 16:20:47 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>QoSHwQueueDelete
i4IfIndex[137]
Jun 15 16:20:47 switch-627e4c last message repeated 7 times
Jun 15 16:20:47 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>QoSHwQueueDelete
i4IfIndex[141]
Jun 15 16:20:47 switch-627e4c last message repeated 7 times
Jun 15 16:20:48 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: ==FsHwSetSpeed
sx_api_port_speed_admin_set = 0
Jun 15 16:20:48 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: ==FsHwGetSpeed
sx_api_port_speed_oper_get = 0
Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS
u4IfIndex[89], u1ConfigOption[6]
Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS
u4IfIndex[33], u1ConfigOption[6]
Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS
u4IfIndex[73], u1ConfigOption[6]
Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS
u4IfIndex[121], u1ConfigOption[6]
Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS
u4IfIndex[133], u1ConfigOption[6]
Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS
u4IfIndex[13], u1ConfigOption[6]
Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS
u4IfIndex[81], u1ConfigOption[6]
Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS
u4IfIndex[117], u1ConfigOption[6]
Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS
u4IfIndex[65], u1ConfigOption[6]
.
.
.
```

Related Commands	
Notes	

8.4 Link Diagnostic Per Port

When debugging a system, it is important to be able to quickly identify the root of a problem. The Diagnostic commands enables an insight into the physical layer components where the user is able to see information such as a cable status (plugged/unplugged) or if Auto-Negotiation has failed.

List of possible output messages:

0—No issue observed
1—Port is close by command (see PAOS)
2—AN no partner detected
3—AN ack not received
4—AN next-page exchange failed
5—KR frame lock not acquired
6—KR link inhibit timeout
7—KR Link partner didn’t set receiver ready
8—KR tuning didn’t completed
9—PCS didn’t acquire block lock
10—PCS didn’t acquire AM lock (NO FEC)
11—PCS didn’t get align_status
12—FC FEC is not locked
13—RS FEC is not locked

14—Remote fault received
15—Bad Signal integrity
16—Compliance code mismatch (protocol mismatch between cable and port)
17—Large number of physical errors (high BER)
18—Port is disabled by Ekey
19—Phase EO failure
20—Stamping of non-NVIDIA Cables/Modules
21—Down by PortInfo MAD
22—Disabled by Verification
23—Calibration failure
24—EDR speed is not allowed due to cable stamping: EDR stamping
25—FDR10 speed is not allowed due to cable stamping: FDR10 stamping
26—Port is closed due to cable stamping: Ethernet_compliace_code_zero
27—Port is closed due to cable stamping: 56GE stamping
28—Port is closed due to cable stamping: non-NVIDIA QSFP28
29—Port is closed due to cable stamping: non-NVIDIA SFP28
30—Port is closed, no backplane enabled speed over backplane channel
31—Port is closed, no passive protocol enabled over passive copper channel
32—Port is closed, no active protocol enabled over active channel
33—Port width is does not match the port speed enabled
34—Local speed degradation
35—Remote speed degradation
36—No Partner detected during force mode.
37—Partial link indication during force mode.
38—AN Failure—FEC mismatch during override
39—AN Failure—No HCD
40—VPI protocol don’t match
41—Port is closed, module can’t be set to the enabled rate
42—Bad SI, cable is configured to non optimal rate
1023—Info not available
MNG FW issues (1024—2047):
1024—Cable is unplugged/powered off
1025—Long Range for non MLNX cable/module .
1026—Bus stuck (I2C Data or clock shorted)
1027—bad/unsupported EEPROM
1028—part number list
1029—unsupported cable.
1030—module temperature shutdown
1031—Shorted cable
1032—Power Budget Exceeded
1033—Management force down the port
1034—Module is disabled by command

8.4.1 Link Diagnostic Commands

8.4.1.1 show interfaces ethernet link-diagnostics

	show interfaces ethernet [<interface>] link-diagnostics Displays a specific Ethernet module/port or all Ethernet ports.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config	
History	3.6.4006	
	3.6.4110	Updated command input
Example	switch (config) # show interfaces ethernet link-diagnostics ---------- Interface Code Status ---------- Eth1/1 1024 Cable is unplugged Eth1/2 1024 Cable is unplugged Eth1/3 1024 Cable is unplugged Eth1/4 1024 Cable is unplugged Eth1/5 1024 Cable is unplugged Eth1/6 1024 Cable is unplugged Eth1/7 1024 Cable is unplugged Eth1/8 1024 Cable is unplugged Eth1/9 1024 Cable is unplugged Eth1/10 1024 Cable is unplugged Eth1/11 1024 Cable is unplugged Eth1/12 1024 Cable is unplugged Eth1/13 1024 Cable is unplugged Eth1/14 1024 Cable is unplugged Eth1/15 1024 Cable is unplugged Eth1/16 1024 Cable is unplugged Eth1/17 1024 Cable is unplugged Eth1/18 1024 Cable is unplugged Eth1/19 1024 Cable is unplugged Eth1/20 1024 Cable is unplugged Eth1/21 1024 Cable is unplugged Eth1/22 1024 Cable is unplugged Eth1/23 1024 Cable is unplugged Eth1/24 1024 Cable is unplugged Eth1/25 1024 Cable is unplugged Eth1/26 1024 Cable is unplugged Eth1/27 1024 Cable is unplugged Eth1/28 1024 Cable is unplugged Eth1/29 1024 Cable is unplugged Eth1/30 1024 Cable is unplugged Eth1/31 0 No issue was observed Eth1/32 0 No issue was observed	
Related Commands		
Notes		

8.5 Signal Degradation Monitoring

A system can monitor the Bit Error Rate (BER) in order to ensure a quality of the link. As long as BER observed by the MACLRH layer is low enough, the rate of packet loss is low enough to allow successful operation of the applications running on top of the network.

The system continuously monitors the link BER and compares it to BER limits, when limits are crossed the system can generate an event indicating that link quality is degraded to the network

operator that can take preemptive actions or even disable the low quality link. The BER configuration threshold for No-FEC and Post-FEC configured for 10^-10 will cause the trap to occur.

When Forward Error Correction (FEC) is enabled a network operator can choose to monitor an amount of corrected errors by using the pre-FEC mode, or the amount of errors which the FEC failed to correct (uncorrectable errors) by using the post-FEC mode, when FEC is used then every error detected by the PHY will be monitored.

When link is disabled the system will keep it in shutdown state until the port is explicitly enabled (Explicitly running “shutdown” and then “no shutdown” commands for that port).

8.5.1 Effective-BER Monitoring

Effective-BER is the BER that the MACLRH/Application layer observe. Errors monitored by the Effective-BER may directly result in a packet drop. For links with no error correction, the Effective BER is the BER received by port, and it is monitored based on the received Phy symbols. For links with FEC, the Effective BER represents the rate of errors that the FEC decoder did not manage to correct and were passed to the MACLRH layer. The Effective BER for FEC links is monitored using the FEC decoder uncorrectable codewords data.

8.5.2 Configuring Signal Degradation Monitoring

1. Enable signal degradation monitoring. Run:

   ```
   switch (config) # interfaces ethernet 1/3 signal-degrade
   ```

   If not indicated, the interface is disabled in case of signal degradation.
2. (Optional) To prevent the interface from shutting down in case of signal degradation, run:

   ```
   switch (config) # interfaces ethernet 1/3 signal-degrade no-shutdown
   ```

   a. (Optional) Enable SNMP notifications on signal degradation events. Run:

      ```
      switch (config) # snmp notify event health-module-status
      ```

      Please refer to “Configuring SNMP Notifications (Traps or Informs)” for a general explanation on how to enable SNMP notifications for specific events.
3. (Optional) Enable email notifications on signal degradation events. Run:

   ```
   switch (config) # email notify event health-module-status
   ```

   Signal degradation snmp event comes only when there is an alarm alert of BER limit cross that is being sent only once. There is no SNMP alarm in case of cross down back to normal threshold, nor in the second time in a row the BER is crossed above again. In order to get another alarm on BER limit cross, it is needed to shutdown the interface and enable it again. Please refer to “Email Notifications” for a general explanation on how to enable email notifications for specific events.

8.5.3 Signal Degradation Monitoring Commands

8.5.3.1 signal-degrade

	signal-degrade [no-shutdown] no signal-degrade [no-shutdown] Enables signal degradation operation per interface. The no form of the command disables signal degradation operation per interface.	
Syntax Description	no-shutdown	Does not shutdown an affected interface
Default	Disabled	
Configuration Mode	config interface ethernet	
History	3.6.4110	
Example	`switch (config interface ethernet 1/1) # signal-degrade`	
Related Commands	show interfaces ethernet signal-degrade	
Notes		

8.5.3.2 show interfaces ethernet signal-degrade

	show interfaces ethernet [<slot>/<port>] signal-degrade Displays signal degradation information.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.4110
Example	

switch (config) # show interfaces ethernet signal-degrade

```
-------------------------------------------------------------------------------------
Interface   Admin state                    Monitoring    Action        FEC type
-------------------------------------------------------------------------------------
Eth1/1      Enabled                        Disabled      Shutdown      no-fec/post-fec
Eth1/2      Enabled                        Disabled      Shutdown      no-fec/post-fec
Eth1/3      Enabled                        Disabled      Shutdown      no-fec/post-fec
Eth1/4      Enabled                        Disabled      Shutdown      no-fec/post-fec
Eth1/5      Enabled                        Disabled      Shutdown      no-fec/post-fec
...
```

Related Commands	
Notes	

8.6 Event Notifications

The OS features a variety of supported events. Events are printed in the system log file and can, optionally, be sent to the system administrator via email, SNMP trap or directly prompted to the terminal.

8.6.1 Supported Event Notifications and MIB Mapping

The following table presents the supported events and maps them to their relevant MIB OID.

Event Name	Event Description	MIB OID	Comments
asic-chip-down	ASIC (chip) down	Mellanox-EFM-MIB: asicChipDown	Not supported
cpu-util-high	CPU utilization has risen too high	Mellanox-EFM-MIB: cpuUtilHigh	N/A
dcbx-ets-port-admin-state-trap	DCBX ETS port admin state trap	MELLANOX-DCB-TRAPS-MIB: mellanoxETSPortAdminStateTrap	N/A
dcbx-ets-port-oper-state-trap	DCBX ETS port oper state trap	MELLANOX-DCB-TRAPS-MIB: mellanoxETSPortOperStateTrap	N/A
dcbx-ets-port-peer-state-trap	DCBX ETS port peer state trap	MELLANOX-DCB-TRAPS-MIB: mellanoxETSPortPeerStateTrap	N/A
dcbx-pfc-module-state-change	DCBX PFC module state change	MELLANOX-DCB-TRAPS-MIB: mellanoxPFCModuleStateTrap	N/A
dcbx-pfc-port-admin-state-trap	DCBX PFC port admin state trap	MELLANOX-DCB-TRAPS-MIB: mellanoxPFCPortAdminStateTrap	N/A
dcbx-pfc-port-oper-state-trap	DCBX PFC port oper state trap	MELLANOX-DCB-TRAPS-MIB: mellanoxPFCPortOperStateTrap	N/A
dcbx-pfc-port-peer-state-trap	DCBX PFC port peer state trap	MELLANOX-DCB-TRAPS-MIB: mellanoxPFCPortPeerStateTrap	N/A
disk-space-low	File system free space has fallen too low	Mellanox-EFM-MIB: diskSpaceLow	N/A
health-module-status	Health module status changed	Mellanox-EFM-MIB: systemHealthStatus	N/A
insufficient-fans	Insufficient amount of fans in system	Mellanox-EFM-MIB: insufficientFans	N/A
insufficient-fans-recover	Insufficient amount of fans in system recovered	Mellanox-EFM-MIB: insufficientFansRecover	N/A
insufficient-power	Insufficient power supply	Mellanox-EFM-MIB: insufficientPower	N/A

Event Name	Event Description	MIB OID	Comments
interface-down	An interface's link state has changed to DOWN	RFC1213: linkdown (SNMPv1)	Supported for Ethernet and management interfaces for 1U and blade systems
interface-up	An interface's link state has changed to UP	RFC1213: linkup (SNMPv1)	Supported for Ethernet and management interfaces for 1U and blade systems
internal-bus-error	Internal bus (I2C) error	Mellanox-EFM-MIB: internalBusError	N/A
liveness-failure	A process in the system is detected as hung	Not implemented	N/A
low-power	Low power supply	Mellanox-EFM-MIB: lowPower	N/A
low-power-recover	Low power supply recover	Mellanox-EFM-MIB: lowPowerRecover	N/A
mstp-new-bridge-root	The bridge become the root bridge root of a MSTI	MELLANOX-MSTP-MIB: mstpRootBridgeChange	N/A
mstp-new-root-port	The root port of a MSTI changed	MELLANOX-MSTP-MIB: mstpRootPortChange	N/A
mstp-topology-change	Port in MSTI become forwarding of blocking	MELLANOX-MSTP-MIB: mstpTopologyChange	N/A
N/A	Reset occurred due to over-heating of ASIC	Mellanox-EFM-MIB: asicOverTempReset	Not supported
new_root	Local bridge became a root bridge	Bridge-MIB: newRoot	N/A
ospf-auth-fail	OSPF authentication failure	OSPF-TRAP-MIB: ospfIfAuthFailure	N/A
ospf-config-error	OSPF config error	OSPF-TRAP-MIB: ospfIfConfigError	N/A
ospf-if-rx-bad-packet	Bad OSPF packet received	OSPF-TRAP-MIB: ospfIfRxBadPacket	N/A
ospf-if-state-change	OSPF interface state change	OSPF-TRAP-MIB: ospfIfStateChange	N/A
ospf-lsdb-approaching-overflow	OSPF LSDB is approaching overflow	OSPF-TRAP-MIB: ospfLsdbApproachingOverflow	Not supported
ospf-lsdb-overflow	OSPF LSDB overflow	OSPF-TRAP-MIB: ospfLsdbOverflow	Not supported
ospf-nbr-state-change	OSPF neighbor state change	OSPF-TRAP-MIB: ospfNbrStateChange	N/A
paging-high	Paging activity has risen too high	N/A	Not supported
process-crash	A process in the system has crashed	Mellanox-EFM-MIB: procCrash	N/A
process-exit	A process in the system unexpectedly exited	Mellanox-EFM-MIB: procUnexpectedExit	N/A

Event Name	Event Description	MIB OID	Comments
send-test	Send a test notification	testTrap	Run the CLI command “snmp-server notify send-test”
snmp-authtrap	An SNMPv3 request has failed authentication	Not implemented	N/A
temperature-too-high	Temperature is too high	Mellanox-EFM-MIB: asicOverTemp	N/A
topology_change	Topology change triggered by a local bridge	Bridge-MIB: topologyChange	N/A
unexpected-shutdown	Unexpected system shutdown	Mellanox-EFM-MIB: unexpectedShutdown	N/A
what-just-happened	Aggregated dropped packets	MELLANOX-WJH-MIB: mellanoxWJHEvent	N/A
xstp-new-root-bridge	The bridge became the root bridge of STI	MELLANOX-XSTP-MIB: mellanoxXstpRootBridgeChange	N/A
xstp-root-port-change	XSTP root port changed	MELLANOX-XSTP-MIB: mellanoxXstpRootPortChange	N/A
xstp-topology-change	Port in pvrst become forwarding of blocking	MELLANOX-XSTP-MIB: mellanoxXstpTopologyChange	N/A

8.6.2 Terminal Notifications

To print events to the terminal, set the events you wish to print to the terminal. Run:

```
switch (config) # logging monitor events notice
```

This command prints system events in the severity “notice” to the screen. For example, in case of interface-down event, the following gets printed to the screen.

```
switch (config) #
Wed Jul 10 11:30:42 2022: Interface 1/17 changed state to DOWN
Wed Jul 10 11:30:43 2022: Interface 1/18 changed state to DOWN
switch (config) #
```

8.6.3 Email Notifications

To configure the OS to send you emails for all configured events and failures:

1. Set your mailhub to the IP address to be your mail client’s server – for example, Microsoft Outlook exchange server.

   ```
   switch (config) # email mailhub <IP address>
   ```

2. Add your email address for notifications. Run:

```
switch (config) # email notify recipient <email address>
```

3. Configure the system to send notifications for a specific event. Run:

```
switch (config) # email notify event <event name>
```

4. Show the list of events for which an email is sent. Run:

```
switch (config) # show email events
Failure events for which emails will be sent:
  process-crash: A process in the system has crashed
  unexpected-shutdown: Unexpected system shutdown

Informational events for which emails will be sent:
  asic-chip-down: ASIC (Chip) Down
  cpu-util-high: CPU utilization has risen too high
  cpu-util-ok: CPU utilization has fallen back to normal levels
  disk-io-high: Disk I/O per second has risen too high
  disk-io-ok: Disk I/O per second has fallen back to acceptable levels
  disk-space-low: Filesystem free space has fallen too low
.
.
.
```

5. Have the system send you a test email. Run:

```
switch (config) # email send-test

The last command should generate the following email:
-----Original Message-----
From: Admin User [mailto:do-not-reply@switch.]
Sent: Sunday, May 01, 2011 11:17 AM
To: <name>
Subject: System event on switch: Test email for event notification

==== System information:
Hostname: switch
Version:  <version> 2011-05-01 14:56:31
          ...
Date:     2011/05/01 08:17:29
Uptime:   17h 8m 28.060s

This is a test email.
==== Done.
```

8.6.4 Command Event Notifications

8.6.4.1 email autosupport enable

	email autosupport enable no email autosupport enable Sends automatic support notifications via email. The no form of the command stops sending automatic support notifications via email.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.2.3000
Example	`switch (config) # email autosupport enable`
Related Commands	
Notes	

8.6.4.2 email autosupport event

	email autosupport event <event> no email autosupport event Specifies for which events to send auto-support notification emails. The no form of the command resets auto-support email security mode to its default.	
Syntax Description	event	• process-crash - a process has crashed • process-exit - a process unexpectedly exited • liveness-failure - a process iss detected as hung • cpu-util-high - CPU utilization has risen too high • cpu-util-ok - CPU utilization has fallen back to normal levels • paging-high - paging activity has risen too high • paging-ok - paging activity has fallen back to normal levels • disk-space-low - filesystem free space has fallen too low • disk-space-ok - filesystem free space is back in the normal range • memusage-high - memory usage has risen too high • memusage-ok - memory usage has fallen back to acceptable levels • netusage-high - network utilization has risen too high • netusage-ok - network utilization has fallen back to acceptable levels • disk-io-high - disk I/O per second has risen too high • disk-io-ok - disk I/O per second has fallen back to acceptable levels • unexpected-cluster-join - node has unexpectedly joined the cluster • unexpected-cluster-leave - node has unexpectedly left the cluster • unexpected-cluster-size - the number of nodes in the cluster is unexpected • unexpected-shutdown - unexpected system shutdown • interface-up - an interface's link state has changed to up • interface-down - an interface's link state has changed to down • user-login - a user has logged into the system • user-logout - a user has logged out of the system • health-module-status - health module status • temperature-too-high - temperature has risen too high • low-power - low power supply • low-power-recover - low power supply recover • insufficient-power - insufficient power supply • power-redundancy-mismatch - power redundancy mismatch • insufficient-fans - insufficient amount of fans in system • insufficient-fans-recover - insufficient amount of fans in system recovered • asic-chip-down - ASIC (chip) down • internal-bus-error - internal bus (I2C) error • internal-link-speed-mismatch - internal links speed mismatch
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Example	`switch (config) # email autosupport event process-crash`	
Related Commands		
Notes		

8.6.4.3 email autosupport ssl mode

	email autosupport ssl mode {none | tls | tls-none} no email autosupport ssl mode Configures type of security to use for auto-support email. The no form of the command resets auto-support email security mode to its default.	
Syntax Description	none	Does not use TLS to secure auto-support email.
	tls	Uses TLS over the default server port to secure auto-support email and does not send an email if TLS fails.
	tls-none	Attempts TLS over the default server port to secure auto-support email, and falls back on plaintext if this fails.
Default	tls-none	
Configuration Mode	config	
History	3.2.3000	
Example	`switch (config) # email autosupport ssl mode tls`	
Related Commands		
Notes		

8.6.4.4 email autosupport ssl cert-verify

	email autosupport ssl cert-verify no email autosupport ssl cert-verify Verifies server certificates. The no form of the command does not verify server certificates.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.2.3000
Example	`switch (config) # email autosupport ssl cert-verify`
Related Commands	
Notes	

8.6.4.5 email autosupport ssl ca-list

	email autosupport ssl ca-list {<ca-list-name> | default_ca_list | none} no email autosupport ssl ca-list Configures supplemental CA certificates for verification of server certificates. The no form of the command removes supplemental CA certificate list.	
Syntax Description	default_ca_list	Default supplemental CA certificate list
	none	No supplemental list (uses built-in list only)
Default	default_ca_list	
Configuration Mode	config	

History	3.2.3000
Example	`switch (config) # email autosupport ssl ca-list default_ca_list`
Related Commands	
Notes	

8.6.4.6 email dead-letter

	email dead-letter {cleanup max-age <duration> | enable} no email dead-letter Configures settings for saving undeliverable emails. The no form of the command disables sending of emails to vendor auto-support upon certain failures.	
Syntax Description	duration	Example: “5d4h3m2s” for 5 days, 4 hours, 3 minutes, 2 seconds
	enable	Saves dead-letter files for undeliverable emails
Default	Save dead letter is enabled The default duration is 14 days	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # email dead-letter enable`	
Related Commands	show email	
Notes		

8.6.4.7 email domain

	email domain <hostname-or-ip-address> no email domain Sets the domain name from which the emails appear to come (provided that the return address is not already fully-qualified). This is used in conjunction with the system hostname to form the full name of the host from which the email appears to come. The no form of the command clears email domain override.	
Syntax Description	hostname-or-ip-address	Hostname or IP address of email domain
Default	No email domain	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # email domain my_domain`	
Related Commands	show emails	
Notes		

8.6.4.8 email mailhub

	email mailhub <hostname-or-ip-address> no email mailhub Sets the mail relay to be used to send notification emails. The no form of the command clears the mail relay to be used to send notification emails.	
Syntax Description	hostname-or-ip-address	Hostname or IP address
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # email mailhub 10.0.8.11`	
Related Commands	show email [events]	
Notes		

8.6.4.9 email autosupport mailhub

	email autosupport mailhub <hostname-or-ip-address> no email autosupport mailhub Sets the mail relay to be used for sending autosupport notification emails. The no form of the command clears the mail relay to be used for sending autosupport notification emails.	
Syntax Description	<hostname-or-ip-address>	The mail hub hostname or IP address
Default	N/A	
Configuration Mode	config	
History	3.7.1000	
Example	`switch (config) # email autosupport mailhub 10.10.10.1`	
Related Commands	show email	
Notes		

8.6.4.10 email autosupport recipient

	email autosupport recipient <email-addr> no email autosupport recipient Sets the recipient for autosupport emails. The no form of the command clears the configured autosupport recipient.	
Syntax Description	email-addr	The autosupport recipient email address
Default	N/A	
Configuration Mode	config	
History	3.7.1000	
Example	`switch (config) # email autosupport recipient user@example.com`	

Related Commands	show email
Notes	

8.6.4.11 email mailhub-port

	email mailhub-port <port number> no email mailhub-port Sets the mail relay port to be used to send notification emails. The no form of the command resets the port to its default.	
Syntax Description	hostname-or-ip-address	Port number
Default	25	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # email mailhub-port 125`	
Related Commands	show email	
Notes		

8.6.4.12 email notify event

	email notify event <event> no email notify event <event> Enables sending email notifications for the specified event type. The no form of the command disables sending email notifications for the specified event type.

Syntax Description	event	Available event names: • process-crash – a process has crashed • process-exit – a process unexpectedly exited • liveness-failure – a process iss detected as hung • cpu-util-high – CPU utilization has risen too high • cpu-util-ok – CPU utilization has fallen back to normal levels • paging-high – paging activity has risen too high • paging-ok – paging activity has fallen back to normal levels • disk-space-low – filesystem free space has fallen too low • disk-space-ok – filesystem free space is back in the normal range • memusage-high – memory usage has risen too high • memusage-ok – memory usage has fallen back to acceptable levels • netusage-high – network utilization has risen too high • netusage-ok – network utilization has fallen back to acceptable levels • disk-io-high – disk I/O per second has risen too high • disk-io-ok – disk I/O per second has fallen back to acceptable levels • unexpected-cluster-join – node has unexpectedly joined the cluster • unexpected-cluster-leave – node has unexpectedly left the cluster • unexpected-cluster-size – the number of nodes in the cluster is unexpected • unexpected-shutdown – unexpected system shutdown • interface-up – an interface's link state has changed to up • interface-down – an interface's link state has changed to down • user-login – a user has logged into the system • user-logout – a user has logged out of the system • health-module-status – health module status • temperature-too-high – temperature has risen too high • low-power – low power supply • low-power-recover – low power supply recover • insufficient-power – insufficient power supply • power-redundancy-mismatch – power redundancy mismatch • insufficient-fans – insufficient amount of fans in system • insufficient-fans-recover – insufficient amount of fans in system recovered • asic-chip-down – ASIC (chip) down • internal-bus-error – internal bus (I2C) error • internal-link-speed-mismatch – internal links speed mismatch
Default	No events are enabled	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # email notify event process-crash`	
Related Commands	email autosupport event show email show email events	
Notes	This does not affect auto-support emails. Auto-support can be disabled overall, but if it is enabled, all auto-support events are sent as emails.	

8.6.4.13 email notify recipient

	email notify recipient <email-addr> [class {info | failure} | detail] no email notify recipient <email-addr> [class {info | failure} | detail] Adds an email address from the list of addresses to which to send email notifications of events. The no form of the command removes an email address from the list of addresses to which to send email notifications of events.	
Syntax Description	email-addr	Email address of intended recipient.
	class	Specifies which types of events are sent to this recipient.
	info	Sends informational events to this recipient.
	failure	Sends failure events to this recipient.
	detail	Sends detailed event emails to this recipient.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # email notify recipient user2@autosupport.mydomain.com`	
Related Commands	show email	
Notes		

8.6.4.14 email return-addr

	email return-addr <username> no email domain Sets the username or fully-qualified return address from which email notifications are sent. • If the string provided contains an "@" character, it is considered to be fully-qualified and used as-is. • Otherwise, it is considered to be just the username, and we append "@<hostname>.<domain>". The default is "do-not-reply", but this can be changed to "admin" or whatnot in case something along the line does not like fictitious addresses. The no form of the command resets this attribute to its default.	
Syntax Description	username	Username
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # email return-addr user1`	
Related Commands	show email	
Notes		

8.6.4.15 email return-host

	email return-host no email return-host Includes the hostname in the return address for emails. The no form of the command does not include the hostname in the return address for emails.
Syntax Description	N/A
Default	No return host
Configuration Mode	config
History	3.1.0000
Example	`switch (config) # no email return-host`
Related Commands	show email
Notes	This only takes effect if the return address does not contain an “@” character

8.6.4.16 email send-test

	email send-test Sends test-email to all configured event and failure recipients.
Syntax Description	N/A
Default	No return host
Configuration Mode	config
History	3.1.0000
Example	`switch (config) # email send-test`
Related Commands	show email [events]
Notes	

8.6.4.17 email ssl mode

	email ssl mode {none | tls | tls-none} no email ssl mode Sets the security mode(s) to try for sending email. The no form of the command resets the email SSL mode to its default.	
Syntax Description	none	No security mode, operates in plaintext
	tls	Attempts to use TLS on the regular mailhub port, with STARTTLS. If this fails, it gives up.
	tls-none	Attempts to use TLS on the regular mailhub port, with STARTTLS. If this fails, it falls back on plaintext.
Default	default-cert	
Configuration Mode	config	
History	3.2.3000	
Example	`switch (config) # email ssl mode tls-none`	

Related Commands	show email
Notes	

8.6.4.18 email ssl cert-verify

	email ssl cert-verify no email ssl cert-verify Enables verification of SSL/TLS server certificates for email. The no form of the command disables verification of SSL/TLS server certificates for email.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.2.3000
Example	`switch (config) # email ssl cert-verify`
Related Commands	show email
Notes	This command has no impact unless TLS is used.

8.6.4.19 email ssl ca-list

	email ssl ca-list {<ca-list-name> | default-ca-list | none} no email ssl ca-list Specifies the list of supplemental certificates of authority (CA) from the certificate configuration database that is to be used for verification of server certificates when sending email using TLS, if any. The no form of the command uses no list of supplemental certificates.	
Syntax Description	ca-list-name	Specifies CA list name
	default-ca-list	Uses default supplemental CA certificate list
	none	Uses no list of supplemental certificates
Default	default-ca-list	
Configuration Mode	config	
History	3.2.3000	
Example	`switch (config) # email ssl ca-list none`	
Related Commands	show email	
Notes	This command has no impact unless TLS is used, and certificate verification is enabled.	

8.6.4.20 show email

	show email Displays email configuration or events for which email should be sent upon.
Syntax Description	N/A

Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config) # show email Mail hub: 10.0.8.70 Mail hub port: 25 Domain override: Return address: do-not-reply Include hostname in return address: yes Current reply address: do-not-reply@<hostname> Security mode: tls-none Verify server cert: yes Supplemental CA list: default-ca-list Dead letter settings: Save dead.letter files: yes Dead letter max age: 14 days Email notification recipients: No recipients configured. Autosupport emails Enabled: no Recipient: Mail hub: Security mode: tls-none Verify server cert: yes Supplemental CA list: default-ca-list
Related Commands	
Notes	

8.6.4.21 show email events

	show email events Displays list of events for which notification emails are sent.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000

Example	switch (config) # show email events
	Failure events for which emails will be sent:
	  expected-shutdown: Expected system shutdown
	  process-crash: A process in the system has crashed
	  unexpected-shutdown: Unexpected system shutdown
	
	Informational events for which emails will be sent:
	  asic-chip-down: ASIC (Chip) Down
	  cpu-util-high: CPU utilization has risen too high
	  cpu-util-ok: CPU utilization has fallen back to normal levels
	  disk-io-high: Disk I/O per second has risen too high
	  disk-io-ok: Disk I/O per second has fallen back to acceptable levels
	  disk-space-low: Filesystem free space has fallen too low
	  disk-space-ok: Filesystem free space is back in the normal range
	  health-module-status: Health module Status
	  insufficient-fans: Insufficient amount of fans in system
	  insufficient-fans-recover: Insufficient amount of fans in system recovered
	  insufficient-power: Insufficient power supply
	  internal-bus-error: Internal bus (I2C) Error
	  internal-link-speed-mismatch: Internal links speed mismatch
	  liveness-failure: A process in the system was detected as hung
	  low-power: Low power supply
	  low-power-recover: Low power supply Recover
	  memusage-high: Memory usage has risen too high
	  memusage-ok: Memory usage has fallen back to acceptable levels
	  netusage-high: Network utilization has risen too high
	  netusage-ok: Network utilization has fallen back to acceptable levels
	  paging-high: Paging activity has risen too high
	  paging-ok: Paging activity has fallen back to normal levels
	  power-redundancy-mismatch: Power redundancy mismatch
	  process-exit: A process in the system unexpectedly exited
	  sm-restart: Subnet Manager restarted for parameter change
	  sm-start: Subnet Manager started
	  sm-stop: Subnet Manager stopped
	  temperature-too-high: Temperature has risen too high
	  unexpected-cluster-join: A node has unexpectedly joined the cluster
	  unexpected-cluster-leave: A node has unexpectedly left the cluster
	  unexpected-cluster-size: The number of nodes in the cluster is unexpected
	
	All events for which autosupport emails will be sent:
	  liveness-failure: A process in the system was detected as hung
	  process-crash: A process in the system has crashed

Related Commands	
Notes	

8.7 Port Mirroring

Port mirroring enables data plane monitoring functionality which allows the user to send an entire traffic stream for testing. Port mirroring sends a copy of packets of a port's traffic stream, called "mirrored port", into an analyzer port. Port mirroring is used for network monitoring. It can be used for intrusion detection, security breaches, latency analysis, capacity and performance matters, and protocol analysis.

The following figure provides an overview of the mirroring functionality.

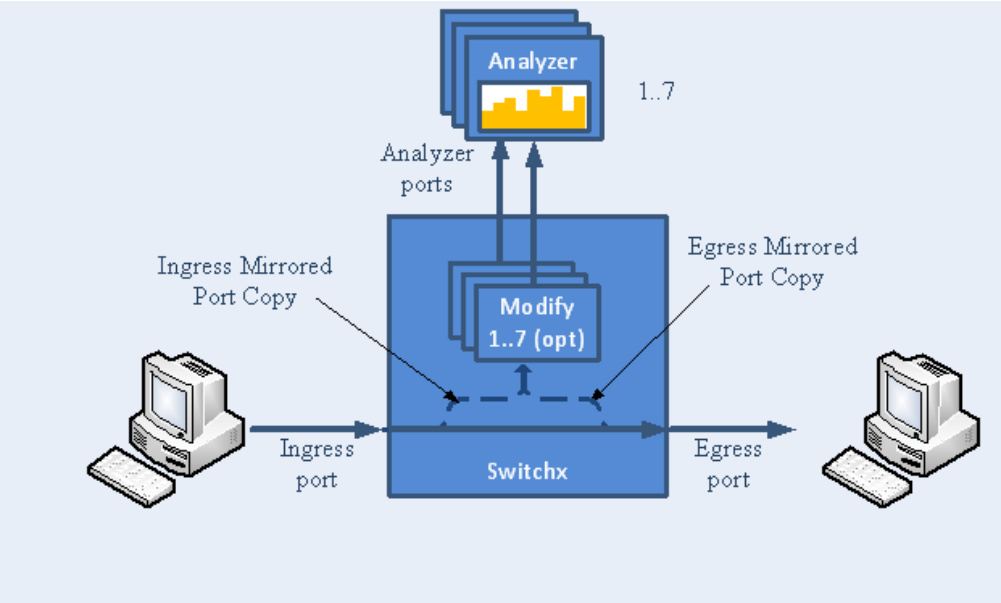

There is no limitation on the number of mirroring sources and more than a single source can be mapped to a single analyzer destination.

8.7.1 Mirroring Sessions

Port mirroring is performed by configuring mirroring sessions. A session is an association of a mirror port (or more) and an analyzer port.

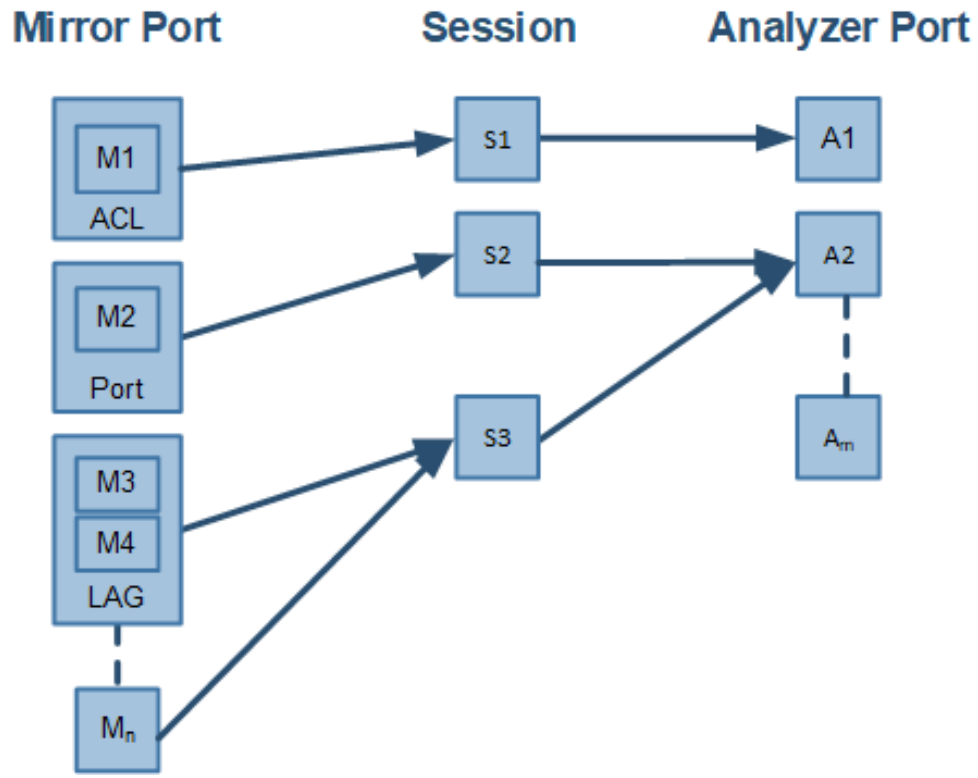

A mirroring session is a monitoring configuration mode that has the following parameters:

Parameter	Description	Access
Source interface(s)	List of source interfaces to be mirrored.	RW
Destination interface	A single analyzer port through which all mirrored traffic egress.	RW
Header format	The format and encapsulation of the mirrored traffic when sent to analyzer.	RW
Truncation	Enabling truncation segments each mirrored packet to 64 bytes.	RW
Congestion control	Controls the behavior of the source port when destination port is congested.	RW
Admin state	Administrative state of the monitoring session.	RW

8.7.1.1 Source Interface

The source interface (mirror port) refers to the interface from which the traffic is monitored. Port mirroring does not affect the switching of the original traffic. The traffic is simply duplicated and sent to the analyzer port. Traffic in any direction (either ingress, egress or both) can be mirrored.

There is no limitation on the number of the source interfaces mapped to a mirroring session.

> Ingress and egress traffic flows of a specific source interface can be mapped to two different sessions.

> There is an option to filter out the specific traffic that needs to be mirrored from the source port by using an ACL action of "monitor session" type. For more details, see "monitor session" command in the ACL Commands section.

8.7.1.1.1 LAG

The source interface can be a physical interface or a LAG.

Port mirroring can be configured on a LAG interface but not on a LAG member. When a port is added to a mirrored LAG it inherits the LAG's mirror configuration. However, if port mirroring configuration is set on a port, that configuration must be removed prior to adding the port to a LAG interface.

When a port is removed from a LAG, the mirror property is switched off for that port.

8.7.1.1.2 Control Protocols

All control protocols captured on the mirror port are forwarded to the analyzer port in addition to their normal treatment. For example LACP, STP, and LLDP are forwarded to the analyzer port in addition to their normal treatment by the CPU.

Exceptions to the behavior above are the packets that are being handled by the MAC layer, such as pause frames.

8.7.1.2 Destination Interface

The destination interface is an analyzer port to which mirrored traffic is directed. The mirrored packets are duplicated, optionally modified, and sent to the analyzer port. Spectrum platforms support up to only 3 analyzer ports, where any mirror port can be mapped to any analyzer port and more than a single mirror port can be mapped to a single analyzer port.

Packets can be forwarded to any destination using the command "destination interface".

The analyzer port supports status and statistics as any other port.

8.7.1.2.1 LAG

The destination interface cannot be a member of LAG when the header format is local.

8.7.1.2.2 Control Protocols

The destination interface may also operate in part as a standard port, receiving and sending out non-mirrored traffic. When the header format is configured as a local port, ingress control protocol packets that are received by the local analyzer port get discarded.

8.7.1.2.3 Advanced MTU Considerations

The analyzer port, like its counterparts, is subject to MTU configuration. It does not send packets longer than configured.

When the analyzer port sends encapsulated traffic, the analyzer traffic has additional headers and therefore longer frame. The MTU must be configured to support the additional length, otherwise, the packet is truncated to the configured MTU.

The system on the receiving end of the analyzer port must be set to handle the egress traffic. If it is not, it might discard it and indicate this in its statistics (packet too long).

8.7.1.3 Header Format

Ingress traffic from the source interface can be manipulated in several ways depending on the network layout using the command header-format.

If the analyzer system is directly connected to the destination interface, then the only parameters that can be configured on the port are the MTU, speed and port based flow control. Priority flow control is not supported is this case. However, if the analyzer system is indirectly connected to the destination interface, there are two options for switching the mirrored data to the analyzer system:

- A VLAN tag may be added to the Ethernet header of the mirrored traffic
- An Ethernet header can be added with include a new destination address and VLAN tag

> It must be taken into account that adding headers increases packet size.

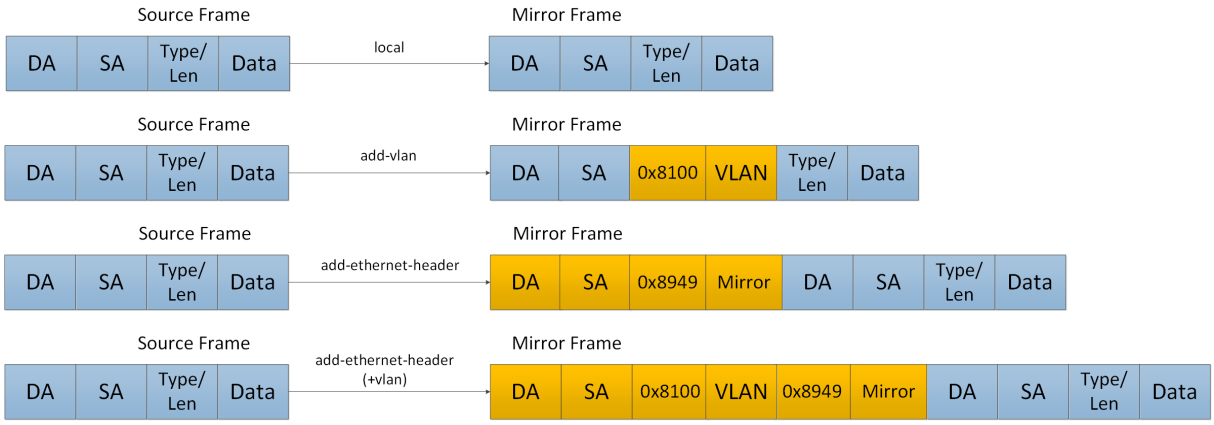

8.7.1.4 Congestion Control

The destination ports might receive pause frames that lead to congestion in the switch port. In addition, too much traffic directed to the analyzer port (for example 40GbE mirror port is directed into 10GbE analyzer port) might also lead to congestion.

In case of congestion:

- When best effort mode is enabled on the analyzer port, Spectrum drops excessive traffic headed to the analyzer port using tail drop mechanism, however, the regular data (mirrored data heading to its original port) does not suffer from a delay or drops due to the analyzer port congestion.
- When the best effort mode on the analyzer port is disabled, the Spectrum does not drop the excessive traffic. This might lead to buffer exhaustion and data path packet loss.

The default behavior in congestion situations is to drop any excessive frames that may clog the system.

> ETS, PFC and FC configurations do not apply to the destination port.

8.7.1.5 Truncation

When enabled, the system can truncate the mirrored packets into smaller 64-byte packets (default) which is enough to capture the packets' L2 and L3 headers.

> The size of the original mirrored packet (before adding the encapsulation headers, and including the 4 bytes frame check sequence (FCs)) is truncated to 64 bytes.

8.7.2 Configuring Mirroring Sessions

The following figure presents two network scenarios with direct and remote connectivity to the analyzer equipment. Direct connectivity is when the analyzer is connected to the analyzer port of the switch. In this case there is no need for adding an L2 header to the mirrored traffic. Remote

connectivity is when the analyzer is indirectly connected to the analyzer port of the switch. In this situation, adding an L2 header may be necessary depending on the network's setup.

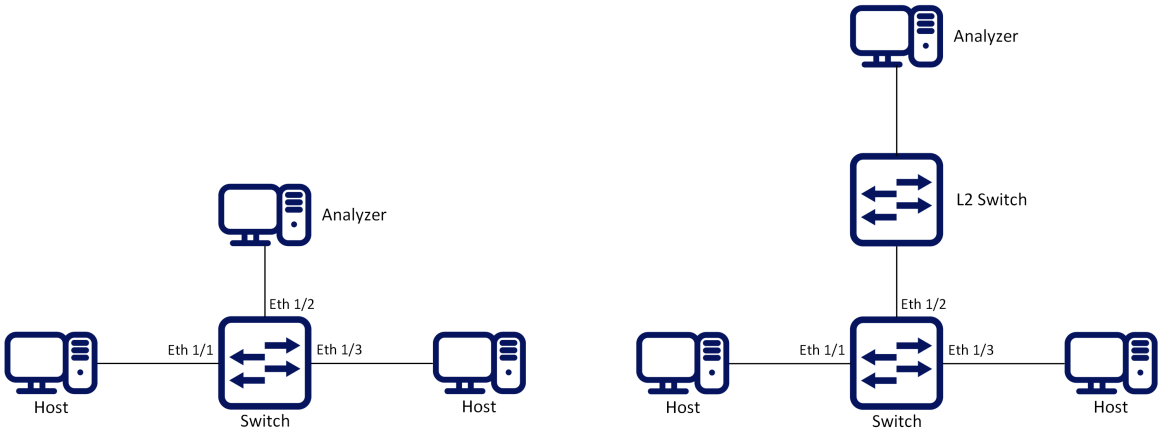

To configure a mirroring session:

1. Create a session. Run:

```
switch (config) # monitor session 1
```

> This command enters a monitor session configuration mode. Upon first implementation the command also creates the session.

2. Add source interface(s). Run:

```
switch (config monitor session 1) # add source interface ethernet 1/1 direction both
```

3. Add destination interface. Run:

```
switch (config monitor session 1) # destination interface ethernet 1/2
```

4. (Optional) Set header format. Run:

```
switch (config monitor session 1) # header-format add-ethernet-header destination-mac 00:0d:ec:f1:a9:c8
add-vlan 10 priority 5
```

> For remote connectivity use the header formats "add-vlan" or "add-ethernet-header". For local connectivity, use "local".

5. (Optional) Truncate the mirrored traffic to 64-byte packets. Run:

```
switch (config monitor session 1) # truncate
```

6. (Optional) Set congestion control. Run:

```
switch (config monitor session 1) # congestion pause-excessive-frames
```

The default for this command is to drop excessive frames. The “pause-excessive-frames” parameter uses flow control to regulate the traffic from the source interfaces.

If the parameter “pause-excessive-frame” is selected, make sure that flow control is enabled on all source interfaces on the ingress direction of the monitoring session using the command “flowcontrol” in the interface configuration mode.

7. Enable the session. Run:

```
switch (config monitor session 1) # no shutdown
```

8.7.3 Verifying Mirroring Sessions

To verify the attributes of a specific mirroring session:

```
switch (config) # show monitor session 1
Session 1:
  Admin:  Enable
  Status: Up
  Truncate: Enable
  Destination interface: eth1/2
  Congestion type: pause-excessive-frames
  Header format: add-ethernet-header
  -switch priority: 5

Source interfaces
--------------------
Interface  Direction
--------------------
eth1/1     both
```

To verify the attributes of running mirroring sessions:

```
switch (config) # show monitor session summary
Flags: i ingress, e egress, b both

------------------------------------------------------------
Session  Admin      Status  Mode       Destination  Source
------------------------------------------------------------
1        Enable     Up      add-eth    eth1/2       eth1/1(b)
2        Disable    Down    add-vlan   eth1/2       eth1/8(i), po1(e)
3        Enable     Up      add-eth    eth1/5       eth1/18(e)
7        Disable    Down    local
```

8.7.4 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community post:

- HowTo Configure Port Mirroring on NVIDIA Ethernet Switches

8.7.5 Port Mirroring Commands

8.7.5.1 monitor session

	monitor session <session-id> no monitor session <session-id> Creates session and enters monitor session configuration mode upon using this command for the first time. The no form of the command deletes the session.	
Syntax Description	session-id	The monitor session ID Range in Spectrum: 1-3 Range in Spectrum-2: 1-8
Default	N/A	
Configuration Mode	config	
History	3.3.3500	
	3.8.1000	Updated syntax
	3.9.1000	Updated notes and "session-id" range
Example	`switch (config)# monitor session 1` `switch (config monitor session 1)#`	
Related Commands	recirculation what-just-happened buffer enable	
Notes	• On Spectrum systems, the maximum number of monitor sessions that can be configured is 2 if a recirculation port is configured, and 3 if not. • On Spectrum-2 systems, the maximum number of monitor sessions that can be configured is 7 if what-just-happened buffer is enabled, and 8 if not.	

8.7.5.2 destination interface

	destination interface <type> <number> [force] no destination interface Sets the egress interface number. The no form of the command deletes the destination interface.	
Syntax Description	interface	Sets the interface type and number (e.g. ethernet 1/2)
	force	Eliminates the need to shutdown the port prior to the operation
Default	no destination interface	
Configuration Mode	config monitor session	
History	3.3.3500	
	3.3.4100	Added force parameter
	3.6.4006	Added note
Example	`switch (config monitor session 1) # destination interface ethernet 1/2`	
Related Commands		

Notes	• Port cannot be used as destination port in monitor session when storm-control is configured on port • Force command cannot remove storm-control configuration. Error output: “Configuration error, storm control is configured on port”. • When removing an interface from a monitor session it gains the default attributes of Ethernet ports

8.7.5.3 shutdown

	shutdown no shutdown Disables the session. The no form of the command enables the session.	
Syntax Description	interface	Sets the interface type and number (e.g. ethernet 1/2)
	force	Eliminates the need to shutdown the port prior to the operation
Default	Disabled	
Configuration Mode	config monitor session	
History	3.3.3500	
	3.3.4100	Added force parameter
	3.6.4006	Added note
Example	`switch (config monitor session 1) # no shutdown`	
Related Commands		
Notes		

8.7.5.4 add source interface direction

	add source interface <type> <number> direction <d-type> no source interface <type> <number> Adds a source interface to the mirrored session. The no form of the command deletes the source interface.	
Syntax Description	interface	Sets the interface type and number (e.g. ethernet 1/2)
	direction	Configures the direction of the mirrored traffic. The options are as follows: • egress - monitors egress traffic • ingress - monitors ingress traffic • both - monitors egress and ingress traffic
Default	N/A	
Configuration Mode	config monitor session	
History	3.3.3500	
Example	`switch (config monitor session 1) # add source interface ethernet 1/1 direction ingress`	
Related Commands		
Notes	• If mirroring is configured in one direction (e.g. ingress) on an interface and then is configured in the other direction (e.g. egress), then the ultimate setting is “both”	

8.7.5.5 header-format

	header-format {local [switch-priority <sp>] | add-vlan <vlan-id> [priority <prio>] [switch-priority <sp>] | add-ethernet-header destination-mac <mac-address> [add-vlan <vlan-id> [priority <prio>]] [switch-priority <sp>]} no header-format Sets the header format of the mirrored traffic. The no form of the command resets the parameter values back to default.	
Syntax Description	local	The mirrored header of the frame is not changed
	switch-priority	Changes the egress switch priority of the frame Range: 0-7
	add-vlan	An 802.1q VLAN tag is added to the frame
	priority	The priority to be added to the Ethernet header Range: 0-7
	add-ethernet-header	Adds an Ethernet header to the mirrored frame
	destination-mac	The destination MAC address of the added Ethernet frame
Default	no-change vlan 1 priority 0 traffic-class 0	
Configuration Mode	config monitor session	
History	3.3.3500	
	3.5.1000	Added switch-priority parameter
	3.8.2000	Updated switch-priority
Example	`switch (config monitor session 1) # header-format add-ethernet-header destination-mac 00:0d:ec:f1:a9:c8 add-vlan 10 priority 5 switch-priority 2`	
Related Commands		
Notes	If add-ethernet-header is used, the source MAC address is the one of the outgoing Ethernet port.	

8.7.5.6 truncate

	truncate no truncate Truncates the mirrored frames to 64-byte packets. The no form of the command disables truncation.	
Syntax Description	N/A	
Default	no truncate	
Configuration Mode	config monitor session	
History	3.3.3500	
	3.9.0500	Added note

Example	`switch (config monitor session 1) # truncate`
Related Commands	
Notes	• This command applies for all sessions on the same analyzer port • The size of the original mirrored packet (before adding the encapsulation headers, and including the 4 bytes frame check sequence (FCs)) is truncated to 64 bytes

8.7.5.7 congestion

	congestion [drop-excessive-frames | pause-excessive-frames] no congestion Sets the system's behavior when congested. The no form of the command disables truncation.	
Syntax Description	drop-excessive-frames	Drops excessive frames
	pause-excessive-frames	Pauses excessive frames
Default	drop-excessive-frames	
Configuration Mode	config monitor session	
History	3.3.3500	
Example	`switch (config monitor session 1) # congestion pause-excessive-frames`	
Related Commands		
Notes	This command applies for all sessions on the same analyzer port	

8.7.5.8 show monitor session

	show monitor session <session-id> Displays monitor session configuration and status.	
Syntax Description	session-id	The monitor session ID Range: 1-7
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.3500	
	3.6.5000	Updated Example

Example	`switch (config) # show monitor session 1`
	`Session 1:`
	`  Admin: Disable`
	`  Status: Down`
	`  Truncate: Disable`
	`  Destination interface: N/A`
	`  Congestion type: drop-excessive-frames`
	`  Header format: local`
	`  -switch priority: 0`
	`Source interfaces`
	`--------------------`
	`Interface  Direction`
	`--------------------`
	`eth1/1  both`	

Related Commands	
Notes	

8.7.5.9 show monitor session summary

	show monitor session summary Displays monitor session configuration and status summary.	
Syntax Description	session-id	The monitor session ID Range: 1-7
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.3500	
	3.6.5000	Updated Example
Example		

```
switch (config) # show monitor session summary
Flags: i ingress, e egress, b both

------------------------------------------------------------
Session  Admin      Status  Mode      Destination  Source
------------------------------------------------------------
1        Disable    Down    local     N/A          eth1/1(b)
2        Disable    Down    add-vlan  eth1/2       eth1/8(i)
```

Related Commands	
Notes	

8.8 sFlow

sFlow (ver. 5) is a procedure for statistical monitoring of traffic in networks.NVIDIA Onyx supports an sFlow sampling mechanism (agent), which includes collecting traffic samples and data from counters. The sFlow datagrams are then sent to a central collector.

The sampling mechanism must ensure that any packet going into the system has an equal chance of being sampled, irrespective of the flow to which it belongs. The sampling mechanism provides the collector with periodical information on the amount (and load) of traffic per interface by loading the counter samples into sFlow datagrams.

The sFlow packets are encapsulated and sent in UDP over IP. The UDP port number that is used is the standard 6343 by default.

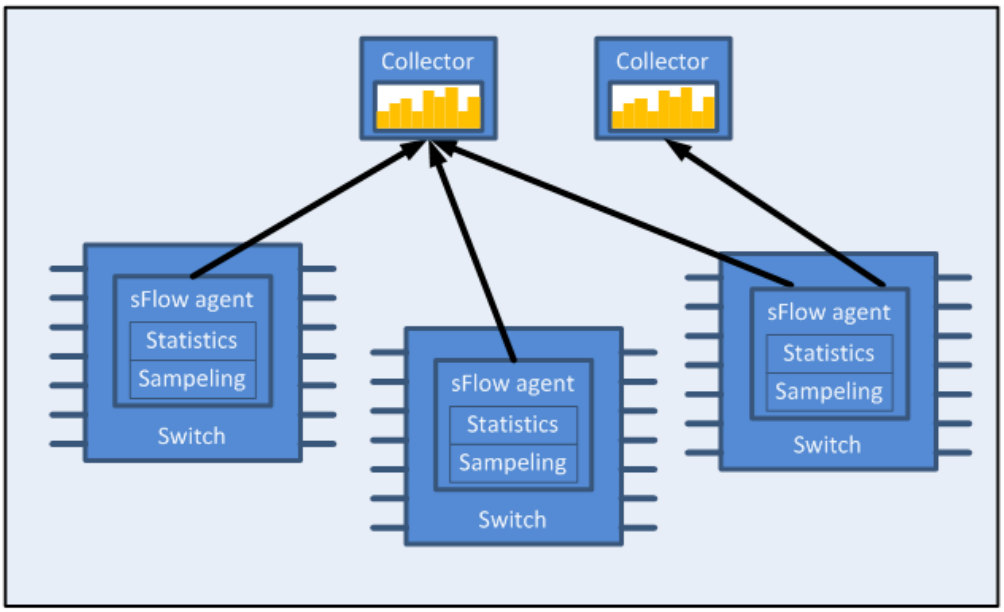

8.8.1 Flow Samples

The sFlow agent samples the data path based on packets.

Truncation and sampling rate are the two parameters that influence the flow samples. In case of congestion the flow samples can be truncated to a predefined size before it is assigned to the CPU. The truncation can be set to any value between 64 to 256 bytes with the default being 128 bytes.

The sampling rate can be adjusted by setting an average rate. The system assures that a random number of packets is sampled, however, the sample rate on average converges to the configured rate. Valid values range between 4000 to 16777215 packets.

The sFlow sends flow samples using an expanded flow sample format.

8.8.2 Statistical Samples

The sFlow agent samples interface counters time based. Polling interval is configurable to any value between 5-3600 seconds with the default being 20 seconds.

The sFlow sends statistical samples using an expanded counter sample format.

The following statistics are gathered by the CPU:

Counter	Description
Total packets	The number of packets that pass through sFlow-enabled ports
Number of flow samples	The number of packets that are captured by the sampling mechanism
Number of statistic samples	The number of statistical samples
Number of discarded samples	The number of samples that were discarded
Number of datagrams	The number of datagrams that were sent to the collector

8.8.3 sFlow Datagrams

The sFlow datagrams contain flow samples and statistical samples.

The sFlow mechanism uses IP protocol, therefore if the packet length is more than the interface MTU, it becomes fragmented by the IP stack. The MTU may also be set manually to anything in the range of 200-9216 bytes. The default is 1400 bytes.

8.8.4 Sampled Interfaces

sFlow must be enabled on physical or LAG interfaces that require sampling. When adding a port to a LAG, sFlow must be disabled on the port. If a port with enabled sFlow is configured to be added to a LAG, the configuration is rejected. Removing a port from a LAG disables sFlow on the port regardless of the LAG's sFlow status.

8.8.5 Configuring sFlow

1. Unlock the sFlow commands.

```
switch (config) # protocol sflow
```

2. Enable sFlow on the system.

```
switch (config) # sflow enable
```

3. Enter sFlow configuration mode.

```
switch (config) # sflow
switch (config sflow) #
```

4. Set the central collector's IP.

```
switch (config sflow) # collector-ip 10.10.10.10
```

5. Set the agent-ip used in the sFlow header.

```
switch (config sflow) # agent-ip 20.20.20.20
```

6. (Optional) Set the sampling rate of the mechanism.

```
switch (config sflow) # sampling-rate 16000
```

This means that one every 16000 packet gets collected for sampling.

7. (Optional) Set the maximum size of the data path sample.

```
switch (config sflow) # max-sample-size 156
```

8. (Optional) Set the frequency in which counters are polled.

```
switch (config sflow) # counter-poll-interval 19
```

9. (Optional) Set the maximum size of the datagrams sent to the central collector.

```
switch (config sflow) # max-datagram-size 1500
```

10. Enable the sFlow agent on the desired interfaces.

```
switch (config interface ethernet 1/1)# sflow enable
switch (config interface port-channel 1)# sflow enable
```

8.8.6 Verifying sFlow

To verify the attributes of the sFlow agent:

```
switch (config)# show sflow
sflow protocol: enabled
sflow: enabled
sampling-rate: 16000
max-sampled-size: 156
counter-poll-interval: 19
max-datagram-size: 1500
collector-ip: 10.10.10.10
collector-port: 6343
agent-ip: 20.20.20.20

ingress ports:
Interfaces:
Ethernet: eth1/1
Port-channel: po1

Statistics:
Total Samples: 2000
Number of flow samples: 1200
Estimated Number of flow discarded: 0
Number of statistic samples: 800
Number of datagrams: 300
```

8.8.7 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community post:

- HowTo Configure sFlow on NVIDIA Switches

8.8.8 sFlow Commands

8.8.8.1 protocol sflow

	protocol sflow no protocol sflow Unhides the sFlow commands. The no form of the command deletes sFlow configuration and hides the sFlow commands.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.3.3500

Example	`switch (config) # protocol sflow`
Related Commands	
Notes	

8.8.8.2 sflow enable (global)

	sflow enable no sflow enable Enables sFlow in the system. The no form of the command disables sFlow without deleting the configuration.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.3.3500
Example	`switch (config) # sflow enable`
Related Commands	
Notes	

8.8.8.3 sflow

	sflow Enters sFlow configuration mode.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.3.3500
Example	`switch (config) # sflow` `switch (config sflow) #`
Related Commands	
Notes	

8.8.8.4 sampling-rate

	sampling-rate <rate> no sampling-rate Configures sFlow sampling ratio. The no form of the command resets this parameter to its default value.	
Syntax Description	rate	Configures the number of packets passed before selecting one for sampling Range: 4000-16777215 “0” disables sampling
Default	N/A	

Configuration Mode	config
History	3.3.3500
Example	`switch (config) # protocol sflow`
Related Commands	
Notes	

8.8.8.5 max-sample-size

	max-sample-size <packet-size> no max-sample-size Configures the maximum size of sampled packets by sFlow. The no form of the command resets the parameter to its default value.	
Syntax Description	packet-size	The sampled packet size Range: 64-256 bytes
Default	N/A	
Configuration Mode	config	
History	3.3.3500	
Example	`switch (config sflow) # max-sample-size 165`	
Related Commands		
Notes	Sampled payload beyond the configured size is discarded	

8.8.8.6 counter-poll-interval

	counter-poll-interval <seconds> no counter-poll-interval Configures the sFlow statistics polling interval. The no form of the command resets the parameter to its default value.	
Syntax Description	seconds	The sFlow statistics polling interval in seconds Range: 5-3600 seconds; “0” disables the statistic polling
Default	20 seconds	
Configuration Mode	config	
History	3.3.3500	
Example	`switch (config sflow) # counter-poll-interval 30`	
Related Commands		
Notes		

8.8.8.7 max-datagram-size

	max-datagram-size <packet-size> no max-datagram-size Configures the maximum sFlow packet size to be sent to the collector. The no form of the command resets the parameter to its default value.

Syntax Description	packet-size	The packet size of the packet being sent to the collector Range: 200-9216 bytes
Default	1400 bytes	
Configuration Mode	config	
History	3.3.3500	
Example	`switch (config sflow) # max-datagram-size 9216`	
Related Commands		
Notes	This packet contains the data sample as well as the statistical counter data	

8.8.8.8 collector-ip

	collector-ip <ip-address> [udp-port <udp-port-number>] no collector-ip [<ip-address> udp-port] Configures the collector's IP. The no form of the command resets the parameters to their default values.	
Syntax Description	ip-address	The collector IP address
	udp-port	Configures the collector UDP port number
Default	ip-address: 0.0.0.0 udf-port-number: 6343	
Configuration Mode	config	
History	3.3.3500	
Example	`switch (config sflow) # collector-ip 10.10.10.10`	
Related Commands		
Notes		

8.8.8.9 agent-ip

	agent-ip {<ip-address> | interface [ethernet <slot/port> | port-channel <channel-group>] | <if-name> | loopback <number> | vlan <id>} no agent-ip Configures the IP address associated with this agent. The no form of the command resets the parameters to their default values.	
Syntax Description	interface	Configures a specific Ethernet/LAG interface's agent IP
	if-name	Interface name (e.g. mgmt0, mgmt1)
	ip-address	The sFlow agent's IP address (i.e. the source IP of the packet)
	loopback	Loopback interface number Range: 1-32
	vlan	Interface VLAN Range: 1-4094
Default	ip-address: 0.0.0.0	
Configuration Mode	config	
History	3.3.3500	

	3.3.5200	Updated “interface” parameters
Example	`switch (config sflow) # agent-ip 20.20.20.20`	
Related Commands		
Notes	The IP address here is used in the sFlow header	

8.8.8.10 clear counters

	clear counters Clears sFlow counters.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.3.3500
Example	`switch (config sflow) # clear counters`
Related Commands	
Notes	

8.8.8.11 sflow enable (interface)

	sflow enable no sflow enable Enables sFlow on this interface. The no form of the command disables sFlow on the interface.	
Syntax Description	N/A	
Default	disable no view-port-channel member	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.3.3500	
	3.3.4500	Added MPO configuration mode
Example	`switch (config interface ethernet 1/1)# sflow enable`	
Related Commands		
Notes		

8.8.8.12 show sflow

	show sflow Displays sFlow configuration and counters.
Syntax Description	N/A
Default	N/A

Configuration Mode	Any command mode	
History	3.3.3500	
	3.6.3004	Updated example
	3.6.6000	Updated example
	3.9.2000	Updated example, adding VRF field
Example	switch (config)# show sflow sflow protocol: enabled sflow: enabled VRF name: mgmt sampling-rate: 16000 max-sample-size: 128 counter-poll-interval: 20 max-datagram-size: 1400 ip-agent: 0.0.0.0 ingress ports: Interfaces: Ethernet eth1/2 eth1/1 Statistics: Total Samples: 0 Number of flow samples: 0 Estimated Number of flow discarded: 0 Number of flow statistics samples: 0 Number of datagrams: 0	
Related Commands		
Notes		

8.9 Buffer Histograms Monitoring

As it is becoming increasingly complex to manage networks, and network administrators need more tools to understand network behavior, it is necessary to provide basic information about network performance, identify network bottlenecks, and provide information for the purposes of network optimization and future planning.

Therefore, network administrators are required to constantly review network port behavior, record port buffer consumption, and identify shortage in buffer resources and record flows which lead to the excessive buffer consumption. NVIDIA Onyx provides the following mechanisms to perform these tasks:

- Sampling (histograms)—a network administrator can enable a sampling of the port buffer occupancy, record occupancy changes over time, and provide information for different levels of buffer occupancy, and amount of time the buffer has been occupied during the observation period.
- Thresholds—thresholds may be enabled per port to record the network time when port buffer occupancy crosses the defined threshold and when buffer occupancy drops below it.
- Flow recording—a record of the most active flows which cause an excessive usage of the port buffers may be kept. Once enabled, the system may identify flow patterns and present a user with a list of flows, based on which a network administrator can rearrange distribution of the data flows in the network and minimize data loss.

8.9.1 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community post:

- Understanding Telemetry Sampling on Spectrum Switches

8.9.2 Buffer Histograms and Thresholds Commands

8.9.2.1 protocol telemetry

	protocol telemetry no protocol telemetry Unhides telemetry config CLIs. The no form of the command hides telemetry config CLIs.
Syntax Description	N/A
Default	Hidden
Configuration Mode	config
History	3.6.3004
Example	`switch (config) # protocol telemetry`
Related Commands	
Notes	

8.9.2.2 telemetry shutdown

	telemetry shutdown no telemetry shutdown Disables the telemetry protocol, threshold detection, and histogram fetching for all sampling enabled interfaces without changing any internal configuration. The no form of the command enables telemetry protocol.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.6.3004
Example	`switch (config) # no telemetry shutdown`
Related Commands	protocol telemetry
Notes	

8.9.2.3 telemetry sampling log

	telemetry sampling log <time> no telemetry sampling log <time> Enables the log interval value (histogram fetching) from device. The no form of the command disables the log interval value.

Syntax Description	time	Input range: 100-60000 (in msec)
Default	1000 millisecond	
Configuration Mode	config	
History	3.6.3004	
Example	`switch (config) # telemetry sampling log 1000`	
Related Commands	protocol telemetry	
Notes		

8.9.2.4 telemetry sampling tc

	telemetry sampling tc <0-7> [mcast | ucast] no telemetry sampling tc <0-7> [mcast | ucast] Enables multicast sampling (histogram fetching) on a traffic class for a particular Ethernet interface. The no form of the command disables multicast sampling on a TC for a particular Ethernet interface.	
Syntax Description	mcast	Multicast traffic
	ucast	Unicast traffic
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.3004	
Example	`switch (config 1/2) # telemetry sampling tc 3 mcast`	
Related Commands		
Notes		

8.9.2.5 telemetry threshold

	telemetry threshold tc <0-7> [ucast | mcast] no telemetry threshold tc <0-7> [ucast | mcast] Enables threshold in hardware for a particular traffic class. The no form of the command disables threshold in hardware for a particular traffic class.	
Syntax Description	ucast	Unicast traffic
	mcast	Multicast traffic
Default	Disabled	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.5000	
Example	`switch (config 1/12) # telemetry threshold tc 0 ucast`	
Related Commands		
Notes		

8.9.2.6 telemetry threshold level

	telemetry threshold level <level> no telemetry threshold level <level> Configures the threshold level in the hardware per port. The no form of the command resets the parameter to its default.		
Syntax Description	level	For Spectrum-based systems: Range: 96-1,000,000 Level is set in bytes and in increments of 96	For Spectrum-2 and Spectrum-3-based systems: Range: 144-1,000,000 Level is set in bytes and in increments of 144
Default	69984		
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel		
History	3.6.5000		
	3.9.0900	Added minimum level value for Spectrum-2 and Spectrum-3	
Example	`switch (config 1/12) # telemetry threshold level 288`		
Related Commands			
Notes			

8.9.2.7 telemetry threshold log

	telemetry threshold log no telemetry threshold log Enables logging of threshold events in syslog. The no form of the command disables logging.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.6.4006
Example	`switch (config) # telemetry threshold log`
Related Commands	
Notes	

8.9.2.8 telemetry threshold syslog

	telemetry threshold syslog <time> no telemetry threshold syslog <time> The command sets threshold events logging rate on per hour basis. The no form of the command sets the logging rate back to default.	
Syntax Description	time	Max rate per hour Range: 1-3600

Default	100
Configuration Mode	config
History	3.6.4006
Example	`switch (config) # telemetry threshold syslog 400`
Related Commands	
Notes	

8.9.2.9 clear telemetry

	clear telemetry {threshold | sampling} [interface <type> <port-id>] [tc <0-7> [ucast | mcast]] Clears telemetry data.	
Syntax Description	type	Possible values: ethernet, port-channel, mlag-port-channel
	tc	Traffic class
	mcast	Multicast traffic
	ucast	Unicast traffic
Default	N/A	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.5000	
Example	`switch (config interface ethernet 1/12) # clear telemetry threshold level 288`	
Related Commands		
Notes		

8.9.2.10 clear telemetry threshold

	clear telemetry threshold [interface <type> <if>] Clears threshold and top talker data.	
Syntax Description	type	Available values:ethernet, port-channel, mlag-port-channel
Default	N/A	
Configuration Mode	config	
History	3.6.6105	
Example	`switch (config) # clear telemetry threshold interface ethernet 1/34-1/36`	
Related Commands		
Notes		

8.9.2.11 stats export csv telemetry

	stats export csv telemetry <slot>/<port>[/<subport>]/<tc>-[mcast | ucast][filename <name>] [after * *] [before * *] Exports histograms collected by stats to a csv file.	
Syntax Description	slot/port	Port number
	subport	Subport number to be used if a port is split
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
	3.9.0500	Updated example
Example	`switch (config) # stats export csv telemetry 1/1/4-ucast after 2020/03/16 10:54:58 before 2020/03/16 11:16:24` `Generated report file: telemetry-20200316-111704.csv`	
Related Commands		
Notes		

8.9.2.12 file stats telemetry delete

	file stats telemetry delete <filename> Deletes the given .csv file created by “stats export” command to user directory.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.3004
Example	`switch (config) # file stats telemetry delete telemetry-20171006-102158.csv`
Related Commands	
Notes	

8.9.2.13 file stats telemetry delete latest

	file stats telemetry delete latest Delete the latest stats telemetry file.
Syntax Description	N/A
Default	N/A
Configuration Mode	Configure terminal
History	3.8.1000
Example	`(config) # file stats telemetry delete latest`
Related Commands	file stats telemetry delete <file_name> file stats telemetry delete all
Notes	

8.9.2.14 file stats telemetry delete all

	file stats telemetry delete all Deletes all stats telemetry files from machine.
Syntax Description	N/A
Default	N/A
Configuration Mode	Configure terminal
History	3.8.1000
Example	`(config) # file stats telemetry delete all`
Related Commands	file stats telemetry delete <file_name> file stats telemetry delete latest
Notes	

8.9.2.15 file stats telemetry upload

	file stats telemetry upload <filename> [vrf <vrf-name>] <upload-url> Uploads .csv file created by “stats export” command to user directory.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.3004
	3.9.2000—Added VRF option

Example	`switch (config) # file stats telemetry upload telemetry-20170119-102715.csv scp://username:password@server//directory`
	
	`Password (if required): ******`

Related Commands	
Notes	

8.9.2.16 file stats telemetry upload latest

	file stats telemetry upload latest [vrf <vrf-name>] <upload-url> Upload the latest stats telemetry file to a remote host.
Syntax Description	file stats telemetry upload latest [vrf <vrf-name>] <upload-url>
Default	N/A
Configuration Mode	Configure terminal
History	3.8.1000
	3.9.2000—Added VRF option
Example	`(config) # file stats telemetry upload latest scp://user:pass@10.135.155.8/tmp`
Related Commands	file stats telemetry upload <file_name> file stats telemetry upload all

Notes	

8.9.2.17 file stats telemetry upload all

	file stats telemetry upload all [vrf <vrf-name>] <upload_url> Upload all stats telemetry files to a remote host.
Syntax Description	vrf-name—VRF to be affected. If "vrf-name" parameter is not specified, "default" VRF will be used.
Default	N/A
Configuration Mode	Configure terminal
History	3.8.1000
	3.9.2000—Added VRF option
Example	`(config) # file stats telemetry upload all scp://user:pass@10.135.155.8/tmp`
Related Commands	file stats telemetry upload <file_name> file stats telemetry upload latest
Notes	

8.9.2.18 show telemetry

	show telemetry Displays the global configuration of telemetry properties.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.4000
Example	

```
Telemetry Status                    : Enabled
H/W Sampling Interval(nsec)         : 512
S/W Sampling Interval(ms)           : 1000
Threshold Logging                   : Disabled
Threshold Logging(rate per hour)    : 100

-------------------------------------------------------------------------------------------------------
Interface           TC                  Sampling            Threshold           Level (bytes)
-------------------------------------------------------------------------------------------------------
Eth1/1              N/A                 Disabled            Disabled            N/A
Eth1/2              N/A                 Disabled            Disabled            N/A
Eth1/3              N/A                 Disabled            Disabled            N/A
Eth1/4              N/A                 Disabled            Disabled            N/A
Eth1/5              N/A                 Disabled            Disabled            N/A
```

Related Commands	
Notes	

8.9.2.19 show telemetry sampling tc mcast

	show telemetry sampling <slot>/<port>[/<subport>] tc <tc_id> mcast Displays fetched multicast histogram details for a given tc_id of the Ethernet interface.	
Syntax Description	slot/port	Ethernet port number
	subport	Ethernet subport number to be used if a port is split
	tc_id	Range: 0-7
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Example		

```
switch (config) # show telemetry sampling 1/2 tc 3 mcast

------------------------------------------------------------------------------------------------------------------------
------------------------------
Telemetry histogram: Eth1/2 traffic-class 3 - mcast
        Time                                  Bin sizes (nsec buffer was occupied in bytes range)
------------------------------------------------------------------------------------------------------------------------
------------------------------
01/16/17                2976<       27552       52128       76704       101280      125856      150432
175008      199584      199584>
04:09:07.79936     1000000000       0           0           0           0           0           0
0           0           0
04:09:08.80096     1000000000       0           0           0           0           0           0
0           0           0
04:09:09.80355     1000000000       0           0           0           0           0           0
0           0           0
04:09:10.80518     1000000000       0           0           0           0           0           0
0           0           0
04:09:11.80682     1000000000       0           0           0           0           0           0
0           0           0
```

Related Commands	
Notes	

8.9.2.20 show telemetry sampling tc mcast last

	show telemetry sampling <slot>/<port>[/<subport>] tc <tc_id> mcast last <num_of_entries> Displays last num of fetched multicast histogram details for the given tc_id of the ethernet interface.	
Syntax Description	slot/port	Ethernet port number
	subport	Ethernet subport number to be used if a port is split
	tc_id	Range: 0-7
	num_of_entries	Range: 0-1000
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Example		

```
switch (config) # show telemetry sampling 1/2 tc 3 mcast last 4
-------------------------------------------------------------------------------------------------------------------
-----------------------------------------
Telemetry histogram: Eth1/2 traffic-class 3 - mcast
        Time                                          Bin sizes (nsec buffer was occupied in bytes
range)
-------------------------------------------------------------------------------------------------------------------
-----------------------------------------
01/16/17                2976<                27552       52128       76704       101280      125856
150432      175008      199584      199584>
04:23:38.28864    1000000000                 0           0           0           0           0           0
0           0           0
04:23:39.28977    1000000000                 0           0           0           0           0           0
0           0           0
04:23:40.29111    1000000000                 0           0           0           0           0           0
0           0           0
04:23:41.29259    1000000000                 0           0           0           0           0           0
0           0           0
```

Related Commands	
Notes	If the requested entries are more than what the DB contains, it prints the amount in the table.

8.9.2.21 show telemetry sampling tc ucast

	show telemetry sampling <slot>/<port>[/<subport>] tc <tc_id> ucast Displays fetched unicast histogram details for a given TC ID of the Ethernet interface.	
Syntax Description	slot/port	Ethernet port number
	subport	Ethernet subport number to be used if a port is split
	tc_id	Range: 0-7
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	

Example

```
switch (config) # show telemetry sampling 1/2 tc 6 ucast
-------------------------------------------------------------------------------------------------------------------
-----------------------------------------
Telemetry histogram: Eth1/2 traffic-class 6 - ucast
        Time                                          Bin sizes (nsec buffer was occupied in bytes
range)
-------------------------------------------------------------------------------------------------------------------
-----------------------------------------
01/13/17                2976<                27552       52128       76704       101280      125856
150432      175008      199584      199584>
08:18:09.67745    1000000000                 0           0           0           0           0           0
0           0           0
08:18:10.67850    1000000000                 0           0           0           0           0           0
0           0           0
08:18:11.67953    1000000000                 0           0           0           0           0           0
0           0           0
```

Related Commands	
Notes	

8.9.2.22 show telemetry sampling tc ucast last

	show telemetry sampling <slot>/<port>[/<subport>] tc <tc_id> ucast last <num_of_entries> Displays last number of fetched unicast histogram details for the given traffic class ID of the Ethernet interface.	
Syntax Description	slot/port	Ethernet port number
	subport	Ethernet subport number to be used if a port is split
	tc_id	Range: 0-7
	num_of_entries	Range: 0-1000
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Example		
Related Commands		
Notes	If the requested entries are more than what the DB contains, it prints the amount in the table.	

8.9.2.23 show telemetry threshold

	show telemetry threshold [interface <type> <port-id>] [tc <0-7> [ucast | mcast]] Displays threshold data for either all interfaces or single interface or per interface per traffic class.	
Syntax Description	type	• ethernet • port-channel • mlag-port-channel
	tc	Traffic class
	mcast	Multicast traffic
	ucast	Unicast traffic
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
	3.6.6105	Updated example

Example

```
 switch (config) # show telemetry threshold 1/10-1/13
-------------------------------------------------------------------------------------------------
Event-id   Date      Time        Port        TC     Level       Duration(100 usec)      Repeated
-------------------------------------------------------------------------------------------------
1          09/21/17  10:11:48    Eth 1/10    0      100         102497.61               1
2          09/21/17  10:12:06    Eth 1/10    3      100         85714.76                1

switch (config) # show telemetry threshold interface port-channel 20 tc 2 mcast
-------------------------------------------------------------------------------------------------
Event-id   Date      Time        Port              TC           Level    Duration(100 usec)    Repeated
-------------------------------------------------------------------------------------------------
1          09/21/17  10:11:48    Po20 (Eth 1/1)    2 (mcast)    100      102497.61             1
2          09/21/17  10:12:06    Po20 (Eth 1/1)    2 (mcast)    100      85714.76              1
```

Related Commands	
Notes	The command supports displaying up to 1000 threshold events. As a result, if more than 1000 thresholds configured in total, some interfaces may not be displayed. Therefore, to query thresholds for a specific interface, please use the command “show telemetry threshold interface <type> <id>”.

8.9.2.24 show files stats telemetry

	show files stats telemetry [filename] Displays all files created by the command “stats export csv telemetry”.	
Syntax Description	filename	Displays stats for the specified file
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
	3.6.8008	Updated example

Example	`switch (config) # show files stats telemetry telemetry-20180527-102715.csv`
	`Hostname :test-switch`
	`Report :telemetry histogram`
	`Time lower bound(UTC) :2018/05/28 05:58:10`
	`Time upper bound(UTC) :2018/05/28 05:58:25`
	`Export time(UTC) :2018/05/28 06:00:06`
	`Time lower bound :2018/05/28 08:58:10 +0300`
	`Time upper bound :2018/05/28 08:58:25 +0300`
	`Export time :2018/05/28 09:00:06 +0300`
	`System version :X86_64 sys_test 2018-05-15 04:02:13 x86_64`	

Related Commands	stats export csv telemetry	
Notes		

8.10 Statistics and Alarms

8.10.1 Commands

8.10.1.1 stats alarm clear

	stats alarm <alarm ID> clear Clears alarm state.

Syntax Description	alarm ID	Alarms supported by the system, for example: • cpu_util_indiv – average CPU utilization too high: percent utilization • disk_io – operating System Disk I/O per second too high: kilobytes per second • fs_mnt – free filesystem space too low: percent of disk space free • intf_util – network utilization too high: bytes per second • memory_pct_used – too much memory in use: percent of physical memory used • paging – paging activity too high: page faults • temperature – temperature is too high: degrees
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # stats alarm cpu_util_indiv clear`	
Related Commands	show stats alarm	
Notes		

8.10.1.2 stats alarm enable

	stats alarm <alarm-id> enable no stats alarm <alarm-id> enable Enables the alarm. The no form of the command disables the alarm, notifications will not be received.	
Syntax Description	alarm ID	Alarms supported by the system, for example: • cpu_util_indiv – average CPU utilization too high: percent utilization • disk_io – operating System Disk I/O per second too high: kilobytes per second • fs_mnt – free filesystem space too low: percent of disk space free • intf_util – network utilization too high: bytes per second • memory_pct_used – too much memory in use: percent of physical memory used • paging – paging activity too high: page faults • temperature – temperature is too high: degrees
Default	The default is different per alarm-id	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # stats alarm cpu_util_indiv enable`	
Related Commands	show stats alarm	
Notes		

8.10.1.3 stats alarm event-repeat

	stats alarm <alarm ID> event-repeat {single | while-not-cleared} no stats alarm <alarm ID> event-repeat Configures repetition of events from this alarm. The no form of this command resets this parameter to its default.	
Syntax Description	alarm ID	Alarms supported by the system, for example: • cpu_util_indiv - average CPU utilization too high: percent utilization • disk_io - operating System Disk I/O per second too high: kilobytes per second • fs_mnt - free filesystem space too low: percent of disk space free • intf_util - network utilization too high: bytes per second • memory_pct_used - too much memory in use: percent of physical memory used • paging - paging activity too high: page faults • temperature - temperature is too high: degrees
	single	Does not repeat events: only sends one event whenever the alarm changes state.
	while-not-cleared	Repeats error events until the alarm clears.
Default	single	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # stats alarm cpu_util_indiv event-repeat single`	
Related Commands	show stats alarm	
Notes		

8.10.1.4 stats alarm {rising | falling}

	stats alarm <alarm ID> {rising | falling} {clear-threshold | error-threshold} <threshold-value> Configure alarms thresholds.	
Syntax Description	alarm ID	Alarms supported by the system, for example: • cpu_util_indiv - average CPU utilization too high: percent utilization • disk_io - operating System Disk I/O per second too high: kilobytes per second • fs_mnt - free filesystem space too low: percent of disk space free • intf_util - network utilization too high: bytes per second • memory_pct_used - too much memory in use: percent of physical memory used • paging - paging activity too high: page faults • temperature - temperature is too high: degrees
	falling	Configures alarm for when the statistic falls too low
	rising	Configures alarm for when the statistic rises too high

	error-threshold	Sets threshold to trigger falling or rising alarm
	clear-threshold	Sets threshold to clear falling or rising alarm
	threshold-value	The desired threshold value, different per alarm
Default	Default is different per alarm-id	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # stats alarm cpu_util_indiv falling clear-threshold 10`	
Related Commands	show stats alarm	
Notes	Not all alarms support all four thresholds.	

8.10.1.5 stats alarm rate-limit

	stats alarm <alarm ID> rate-limit {count <count-type> <count> | reset | window <window-type> <duration>} Configures alarms rate limit.	
Syntax Description	alarm ID	Alarms supported by the system, for example: • cpu_util_indiv – average CPU utilization too high: percent utilization • disk_io – operating System Disk I/O per second too high: kilobytes per second • fs_mnt – free filesystem space too low: percent of disk space free • intf_util – network utilization too high: bytes per second • memory_pct_used – too much memory in use: percent of physical memory used • paging – paging activity too high: page faults • temperature – temperature is too high: degrees
	count-type	Long medium, or short count (number of alarms)
	reset	Set the count and window durations to default values for this alarm
	window-type	Long medium, or short count, in seconds
Default	Short window: 5 alarms in 1 hour Medium window: 20 alarms in 1 day Long window: 50 alarms in 7 days	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # stats alarm paging rate-limit window long 2000`	
Related Commands	show stats alarm	
Notes		

8.10.1.6 stats chd clear

	stats chd <CHD ID> clear Clears CHD counters.	
Syntax Description	CHD ID	CHD supported by the system, for example: • cpu_util – CPU utilization: percentage of time spent • cpu_util_ave – CPU utilization average: percentage of time spent • cpu_util_day – CPU utilization average: percentage of time spent • disk_device_io_hour – storage device I/O read/write statistics for the last hour: bytes • disk_io – operating system aggregate disk I/O average (KB/sec) • fs_mnt_day – filesystem system usage average: bytes • fs_mnt_month – filesystem system usage average: bytes • fs_mnt_week – filesystem system usage average: bytes • intf_day – network interface statistics aggregation: bytes • intf_hour – network interface statistics (same as “interface” sample) • intf_util – aggregate network utilization across all interfaces • memory_day – average physical memory usage: bytes • memory_pct – average physical memory usage • paging – paging activity: page faults • paging_day – paging activity: page faults • eth_day • eth_hour • eth_ip_day • eth_ip_hour
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # stats chd memory_day clear`	
Related Commands	show stats chd	
Notes		

8.10.1.7 stats chd enable

	stats chd <chd-id> enable no stats chd <chd-id> enable Enables the CHD. The no form of the command disables the CHD.

Syntax Description	chd-id	CHD supported by the system, for example: • cpu_util - CPU utilization: percentage of time spent • cpu_util_ave - CPU utilization average: percentage of time spent • cpu_util_day - CPU utilization average: percentage of time spent • disk_device_io_hour - storage device I/O read/write statistics for the last hour: bytes • disk_io - operating system aggregate disk I/O average: KB/sec • fs_mnt_day - filesystem system usage average: bytes • fs_mnt_month - filesystem system usage average: bytes • fs_mnt_week - filesystem system usage average: bytes • intf_day - network interface statistics aggregation: bytes • intf_hour - network interface statistics (same as "interface" sample) • intf_util - aggregate network utilization across all interfaces • memory_day - average physical memory usage: bytes • memory_pct - average physical memory usage • paging - paging activity: page faults • paging_day - paging activity: page faults • eth_day • eth_hour
Default	Enabled	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # stats chd memory_day enable`	
Related Commands	show stats chd	
Notes		

8.10.1.8 stats chd compute time

	stats chd <CHD ID> compute time {interval | range} <number of seconds> Sets parameters for when this CHD is computed.

Syntax Description	CHD ID	Possible IDs: • cpu_util - CPU utilization: percentage of time spent • cpu_util_ave - CPU utilization average: percentage of time spent • cpu_util_day - CPU utilization average: percentage of time spent • disk_device_io_hour - storage device I/O read/write statistics for the last hour: bytes • disk_io - operating system aggregate disk I/O average: KB/sec • fs_mnt_day - filesystem system usage average: bytes • fs_mnt_month - filesystem system usage average: bytes • fs_mnt_week - filesystem system usage average: bytes • intf_day - network interface statistics aggregation: bytes • intf_hour - network interface statistics (same as “interface” sample) • intf_util - aggregate network utilization across all interfaces • memory_day - average physical memory usage: bytes • memory_pct - average physical memory usage • paging - paging activity: page faults • paging_day - paging activity: page faults • eth_day • eth_hour
	interval	Specifies calculation interval (how often to do a new calculation) in number of seconds
	range	Specifies calculation range, in number of seconds
	number of seconds	Number of seconds
Default	Different per CHD	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # stats chd memory_day compute time interval 120`	
Related Commands	show stats chd	
Notes		

8.10.1.9 stats export

	stats export <format> <sample-id> Exports collected information to a file.Can export extended "interface-ethernet", "interface-port-channel", "interface-mlag-port-channel"& "power" samples.	
Syntax Description	memory	Memory utilization
	paging	Paging I/O
	telemetry	Telemetry histogram
	cpu_util	CPU utilization
	power	Power
Default	N/A	
Configuration Mode	config	

History	3.7.1102 3.10.1000: Updated syntax description options
Example	`switch (config) # stats export csv memory`
Related Commands	show stats sample
Notes	

8.10.1.10 stats sample clear

	stats sample <sample ID> clear Clears sample history.	
Syntax Description	sample ID	Possible sample IDs are: • congested • cpu_util – CPU utilization: milliseconds of time spent • disk_device_io – storage device I/O statistics • disk_io – operating system aggregate disk I/O: KB/sec • fan - Fan speed • fs_mnt_bytes – filesystem usage: bytes • fs_mnt_inodes – filesystem usage: inodes • interface – network interface statistics • intf_util – network interface utilization: bytes • memory – system memory utilization: bytes • paging – paging activity: page faults • power – power supply usage • power-consumption • temperature – modules temperature • interface-ethernet – Ethernet counters statistics: counter units • interface-mlag-port-channel – MLAG counters statistics: counter units • interface-port-channel – LAG counters statistics: counter units • eth • eth-abs • eth_ip
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # stats sample temperature clear`	
Related Commands	show stats sample	
Notes		

8.10.1.11 stats sample enable

	stats sample <sample-id> enable no states sample <sample-id> enable Enables the sample. The no form of the command disables the sample.

Syntax Description	sample-id	Possible sample IDs are: • congested • cpu_util - CPU utilization: milliseconds of time spent • disk_device_io - storage device I/O statistics • disk_io - operating system aggregate disk I/O: KB/sec • fan - fan speed • fs_mnt_bytes - filesystem usage: bytes • fs_mnt_inodes - filesystem usage: inodes • interface - network interface statistics • intf_util - network interface utilization: bytes • memory - system memory utilization: bytes • paging - paging activity: page faults • power - power supply usage • power-consumption • temperature - modules temperature • interface-ethernet - Ethernet counters statistics: counter units • interface-mlag-port-channel - MLAG counters statistics: counter units • interface-port-channel - LAG counters statistics: counter units • eth • eth-abs • eth_ip
Default	Enabled	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # stats sample temperature enable`	
Related Commands	show stats sample	
Notes		

8.10.1.12 stats sample interval

	stats sample <sample-id> interval [<interval>] no stats sample <sample-id> interval [<interval>] Sets the sampling interval between taking of sample records. The no form of the command sets interval to default value.

Syntax Description	sample-id	Sample name for which report file should be generated. • congested • cpu_util - CPU utilization: milliseconds of time spent • disk_device_io - storage device I/O statistics • disk_io - operating system aggregate disk I/O: KB/sec • fan - fan speed • fs_mnt_bytes - filesystem usage: bytes • fs_mnt_inodes - filesystem usage: inodes • interface - network interface statistics • intf_util - network interface utilization: bytes • memory - system memory utilization: bytes • paging - paging activity: page faults • power - power supply usage • power-consumption • temperature - modules temperature • interface-ethernet - Ethernet counters statistics: counter units • interface-mlag-port-channel - MLAG counters statistics: counter units • interface-port-channel - LAG counters statistics: counter units • eth • eth-abs • eth_ip
	interval	Measured in seconds. Range: 1 - 86400 (24 hours)
Default	Default for “interface” samples is 60 seconds	
Configuration Mode	config	
History	3.7.1102	
Example	`switch (config) # stats sample interface-ethernet interval 1`	
Related Commands	show stats sample	
Notes		

8.10.1.13 stats sample max-entries

	stats sample <sample-id> max-entries [<max-entries>] no stats sample <sample-id> max-entries [<max-entries>] Sets number of records to be kept in memory for the counter. The no form of the command resets the value to its default.

Syntax Description	sample-id	Sample name for which report file should be generated. • congested • cpu_util - CPU utilization: milliseconds of time spent • disk_device_io - storage device I/O statistics • disk_io - operating system aggregate disk I/O: KB/sec • fan - fan speed • fs_mnt_bytes - filesystem usage: bytes • fs_mnt_inodes - filesystem usage: inodes • interface - network interface statistics • intf_util - network interface utilization: bytes • memory - system memory utilization: bytes • paging - paging activity: page faults • power - power supply usage • power-consumption • temperature - modules temperature • interface-ethernet - Ethernet counters statistics: counter units • interface-mlag-port-channel - MLAG counters statistics: counter units • interface-port-channel - LAG counters statistics: counter units • eth • eth-abs • eth_ip
	max-entries	Number of records Range: 1-1000
Default	Default “interface” samples is 100 records	
Configuration Mode	config	
History	3.7.1102	
Example	`switch (config) # stats sample interface-ethernet max-entries 1000`	
Related Commands	show stats sample	
Notes	• Setting a new value will delete all sample history. • History does not persist after reboot.	

8.10.1.14 stats clear-all

	stats clear-all Clears data for all samples, CHDs, and status for all alarms.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Example	`switch (config) # stats clear-all`
Related Commands	show stats sample
Notes	

8.10.1.15 show stats alarm

	show stats alarm [<alarm-id> [rate-limit]] Displays status of all alarms or the specified alarm.	
Syntax Description	alarm-id	Available values: • cpu_util_indiv – average CPU utilization too high: percent utilization • disk_io – operating System Disk I/O per second too high: kilobytes per second • fs_mnt – free filesystem space too low: percent of disk space free • intf_util – network utilization too high: bytes per second • memory_pct_used – too much memory in use: percent of physical memory used • paging – paging activity too high: page faults • temperature – temperature is too high: degrees
	rate-limit	Displays rate limit parameters.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	

Example	`switch (config) # show stats alarm`
	`Alarm cpu_util_indiv (Average CPU utilization too high):  ok`
	`Alarm disk_io (Operating System Disk I/O per second too high): (disabled)`
	`Alarm fs_mnt (Free filesystem space too low):  ok`
	`Alarm intf_util (Network utilization too high):  (disabled)`
	`Alarm memory_pct_used (Too much memory in use):  (disabled)`
	`Alarm paging (Paging activity too high):  ok`
	`Alarm temperature (Temperature is too high):  ok`	

Related Commands	stats alarm	
Notes		

8.10.1.16 show stats chd

	show stats chd [<chd-id>] Displays configuration of all statistics CHDs.	
Syntax Description	chd-id	Available values: • cpu_util_indiv – average CPU utilization too high: percent utilization • disk_io – operating System Disk I/O per second too high: kilobytes per second • fs_mnt – free filesystem space too low: percent of disk space free • intf_util – network utilization too high: bytes per second • memory_pct_used – too much memory in use: percent of physical memory used • paging – paging activity too high: page faults • temperature – temperature is too high: degrees
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	

Example	switch (config) # show stats chd disk_device_io_hour CHD "disk_device_io_hour" (Storage device I/O read/write statistics for the last hour: bytes): Enabled: yes Source dataset: sample "disk_device_io" Computation basis: data points Interval: 1 data point(s) Range: 1 data point(s)
Related Commands	stats chd
Notes	

8.10.1.17 show stats cpu

	show stats cpu Displays some basic stats about CPU utilization: • the current level • the peak over the past hour • the average over the past hour
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config) # show stats cpu CPU 0 Utilization: 6% Peak Utilization Last Hour: 16% at 2012/02/28 08:47:32 Avg. Utilization Last Hour: 8%
Related Commands	
Notes	

8.10.1.18 show stats sample

	show stats sample [<sample-id>] Displays sampling interval for all samples, or the specified one.

Syntax Description	sample-id	Sample name for which report file should be generated. • congested • cpu_util - CPU utilization: milliseconds of time spent • disk_device_io - storage device I/O statistics • disk_io - operating system aggregate disk I/O: KB/sec • fan - fan speed • fs_mnt_bytes - filesystem usage: bytes • fs_mnt_inodes - filesystem usage: inodes • interface - network interface statistics • intf_util - network interface utilization: bytes • memory - system memory utilization: bytes • paging - paging activity: page faults • power - power supply usage • power-consumption • temperature - modules temperature • interface-ethernet - Ethernet counters statistics: counter units • interface-mlag-port-channel - MLAG counters statistics: counter units • interface-port-channel - LAG counters statistics: counter units • eth • eth-abs • eth_ip
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	

Example	`switch (config) # show stats sample fan`
	`Sample "fan" (Fan speed):`
	`  Enabled:  yes`
	`  Sampling interval: 1 minute 11 seconds`	

Related Commands		
Notes		

8.10.1.19 show stats sample data

	show stats sample <sample-id> data [interface {ethernet | port-channel | mlag-port-channel} <device/port> [counter <counter-name>]] [group name <group-name> [counter <counter-name>]] [max-samples {<max-samples> | all}] Displays history of counter values (i.e., collected information for a sample).

Syntax Description	sample-id	Sample name for which report file should be generated. • congested • cpu_util - CPU utilization: milliseconds of time spent • disk_device_io - storage device I/O statistics • disk_io - operating system aggregate disk I/O: KB/sec • fan - fan speed • fs_mnt_bytes - filesystem usage: bytes • fs_mnt_inodes - filesystem usage: inodes • interface - network interface statistics • intf_util - network interface utilization: bytes • memory - system memory utilization: bytes • paging - paging activity: page faults • power - power supply usage • power-consumption • temperature - modules temperature • interface-ethernet - Ethernet counters statistics: counter units • interface-mlag-port-channel - MLAG counters statistics: counter units • interface-port-channel - LAG counters statistics: counter units • eth • eth-abs • eth_ip
	interface	Allows limiting output to a particular interface's counters
	group	Allows limiting output to a particular group of counters
	counter	Allows limiting output to a particular counter. This option is available only if the option interface or group is chosen.
	max-samples	Allows choosing a number of counter records to display. Range: 1-1000 records. The "all" option is meant for all available records. By default, 20 counter records are displayed.
Default	N/A	
Configuration Mode	Any command mode	
History	3.7.1102 3.8.1000: Modified configuration mode & example 3.9.2000: Modified note and example	
Example		

```
switch (config) # show stats sample interface-ethernet data interface ethernet 1/1 max-samples 1

Sampling data for Interface ethernet counters:
Eth1/1:
    ------------------------------------------------------------------
    Name                              Timestamp              Value
    ------------------------------------------------------------------
    Rx_packets                        2000/12/25 10:27:53     0
    Rx_unicast_packets                2000/12/25 10:27:53     0
    Rx_multicast_packets              2000/12/25 10:27:53     0
    Rx_broadcast_packets              2000/12/25 10:27:53     0
    Rx_bytes                          2000/12/25 10:27:53     0
    Rx_discard_packets                2000/12/25 10:27:53     0
    Rx_error_packets                  2000/12/25 10:27:53     0
    Rx_fcs_errors                     2000/12/25 10:27:53     0
    Rx_undersize_packets              2000/12/25 10:27:53     0
    Rx_oversize_packets               2000/12/25 10:27:53     0
    Rx_pause_packets                  2000/12/25 10:27:53     0
    Rx_unknown_control_opcode         2000/12/25 10:27:53     0
    Rx_symbol_errors                  2000/12/25 10:27:53     0
    Rx_packets_of_64_bytes            2000/12/25 10:27:53     0
    Rx_packets_of_65-127_bytes        2000/12/25 10:27:53     0
    Rx_packets_of_128-255_bytes       2000/12/25 10:27:53     0
    Rx_packets_of_256-511_bytes       2000/12/25 10:27:53     0
    Rx_packets_of_512-1023_bytes      2000/12/25 10:27:53     0
    Rx_packets_of_1024-1518_bytes     2000/12/25 10:27:53     0
    Rx_packets_Jumbo                  2000/12/25 10:27:53     0
    Tx_packets                        2000/12/25 10:27:53     0
    Tx_unicast_packets                2000/12/25 10:27:53     0
    Tx_multicast_packets              2000/12/25 10:27:53     0
    Tx_broadcast_packets              2000/12/25 10:27:53     0
    Tx_bytes                          2000/12/25 10:27:53     0
    Tx_discard_packets                2000/12/25 10:27:53     0
    Tx_error_packets                  2000/12/25 10:27:53     0
    Tx_hoq_discard_packets            2000/12/25 10:27:53     0
    Tx_pause_packets                  2000/12/25 10:27:53     0
    Tx_pause_duration                 2000/12/25 10:27:53     0
```

Related Commands	

Notes	• Filtering keyword depends on chosen <sample-id>. For convenience, “interface” samples such as “interface-ethernet”, “interface-port-channel” and “interface-mlag-port-channel” have interface related keywords for choosing a counters group.
	• Notice that this is a history of counters. Autocompletion and output can contain information for groups (interfaces) that is not present anymore in the system, and vice versa. If counters are not sampled, they will not appear in the output.
	• Output of collected information is implemented only for the following samples:
	  ◦ memory
	  ◦ paging
	  ◦ power
	• interface-port-channel
	• interface-ethernet
	• interface-mlag-port-channel

8.11 Management Information Bases (MIBs)

8.11.1 Calculating of entPhysicalIndex in the Entity MIB

The inventory in the switch system can be accessed through a MIB browser. These devices are indexed (entPhysicalIndex) using three layers:

1. Module layer—includes modules located on system (e.g., cables, fan, power supply, and so forth). See the module type breakdown table for more details.
2. Device layer—a number identifying the specific device that is associated with the module (e.g., ASIC on a leaf, fan on the management board, and so forth).
3. Sensor layer—a number identifying the specific sensor that is associated with the device (e.g., fan sensors, temperature sensors, power sensors, and so forth).

Each layer is assigned a fixed position in the SNMP index number that represent it.

The physical entities in the system (other than port modules) use the following index schema:

Mod. Type ID	Module Index		Device Identifier				Sensor Type and Index	
1	2	3	4	5	6	7	8	9
Layer 1			Layer 2				Layer 3	

Spectrum-2 systems and above use the following index schema for port modules and port module sensors:

Mod. Type ID	Port Module Identifier							Port module Sensor index TX sensors in range 1..39 RX sensors in range 41..79	
1	2	3	4	5	6	7	8	9	10
Layer 1	Layer 2							Layer 3	

Spectrum systems use the following index schema for port modules and port module sensors:

Mod. Type ID	Port Module Identifier					Port Module Sensor Type 0 for TX 1 for RX	Sensor index	
1	2	3	4	5	6	7	8	9
Layer 1	Layer 2					Layer 3		

Module type breakdown:

Number	Description
1	Chassis
2	Management
3	Spine
4	Leaf
5	Fan
6	Power supply

Number	Description
7	BBU
8	x86 CPU
9	Port module
Physical entities—10 digits representation	
1	Port module

Port module 9 digits representation is kept for backwards compatibility.

8.11.2 Examples

- entPhysicalIndex with value 401191311
  - 9 digits representation.
  - Layer 1 is “401”—“4” indicates a leaf (see module type breakdown table) and “01” indicates leaf at index #1 (i.e., leaf 01)
  - Layer 2 is “1913”—this is the identifier for one of the QSFP-ASIC in the system
  - Layer 3 is “11”—this is the identifier for temperature sensor #1
  - The description for this physical entity (appears in entPhysicalDescr column of the MIB) would be: L01/QSFP-ASIC-1/T1
- entPhysicalIndex with value 501020021
  - 9 digits representation.
  - Layer 1 is “501”—“5” indicates a fan (see module type breakdown table) and “01” indicates fan at index #1 (i.e., fan 01)
  - Layer 2 is “0200”—this is the identifier for general fan in the system
  - Layer 3 is “21”—this is the identifier for fan sensor #1
  - The description for this physical entity (appears in entPhysicalDescr column of the MIB) would be: FAN1/FAN/F1
- For entPhysicalIndex with value 1000012700
  - 10 digits representation.
  - Layer 1 is “1”—port module (see module type breakdown table).
  - Layer 2 is “127”—port identifier
  - Layer 3 is “00”—no sensors for this port module
- For entPhysicalIndex with value 1000012742
  - 10 digits representation.
  - Layer 1 is “1”—port module (see module type breakdown table).
  - Layer 2 is “127”—port identifier
  - Layer 3 is “42”—sensor in the range 41..79 indicts an RX sensor

9 Automation Tools

Deploying, provisioning, operating and configuring data center networks is still a largely manual and time-consuming process that is susceptible to human error. Its automation greatly enhances agility, accelerates deployment, increases reliability and improves the performance of critical business applications, and at the bottom line it saves on operational expenditure.

The datacenter is an ecosystem composed of computer servers and storage and networking equipment, while each of these components is managed by a separate team using separate tools. Nowadays it is possible to increase efficiency by allowing IT departments to break down barriers, automate processes and better divide resources across the entire datacenter. Network automation enables IT departments to be more responsive to various, real-time business requirements, and more service-centric in their approach to delivering value.

Additionally, it enables a more efficient method to easily change server configuration and apply it to all affected elements of the infrastructure (e.g. when a new virtual machine is spun up, its corresponding VLAN should be configured automatically).

The transition to automated operation is vital to the data center in each of the following aspects:

- Provisioning and deployment: Instead of a time-consuming manual staging process, new switches enable automatic downloading of the correct image and configuration as soon as they are installed on the rack and booted, automating set-up, configuration and the provisioning process.
- Management and operations: Once the network is up and running, adjustments can be programmed to occur automatically, using analytics to deliver current, consistent and accurate information.
- Orchestration: The network must be synched with all other elements of the data center. When a server or storage configuration is changed, it often requires corresponding changes in the network, which need to take place immediately and automatically.
  To enable data center orchestration, switches should:
  - Support orchestration tools such as OpenStack and CloudStack
  - Support SDN solutions from a variety of vendors, such as Juniper's Contrail Networking product
  - Support IT automation solutions, such as Chef, so the network can be managed in concert with the overall data center infrastructure

The below sections provide detailed guideline on how to use two of the main automation tools (Ansible and SALT stack), enabling higher automation in the data center.

9.1 Ansible

> The final version to support Ansible is 3.9.3220.

9.2 SALT

Salt is a different approach to infrastructure management, founded on the idea that high-speed communication with large numbers of systems can open new capabilities. This approach makes Salt a powerful multitasking system that can solve many specific problems in an infrastructure.

The backbone of Salt is the remote execution engine, which creates a high-speed, secure and bi-directional communication net for groups of systems. On top of this communication system, Salt provides an extremely fast, flexible, and easy-to-use configuration management system called Salt States.

For a list of Salt's Napalm supported modules, please refer to the NAPALM-Onyx github repository.

9.2.1 Installing SaltStack on CentOS 7

1. Install Salt packages:

```
curl -L https://bootstrap.saltstack.com -o install_salt.sh
sudo sh install_salt.sh -P -M
yum install -y salt-master salt-minion salt-ssh salt-syndic salt-cloud salt-api
```

2. Install the Napalm library.

```
yum install epel-release
yum install -y python-pip
yum install libxml2-devel libxslt-devel zlib-devel gcc openssl-devel libffi-devel python-devel
pip install pyzmq --install-option="--zmq=bundled"
pip install napalm
```

9.2.2 Configuring Salt

1. Open the /etc/salt/master file.
2. Replace #interface: 0.0.0.0 with interface: <machine_ip>.
3. Replace #hash_type: md5 with hash_type: sha256.
4. Find file_roots and pillar_rootsand and add the following lines below them:

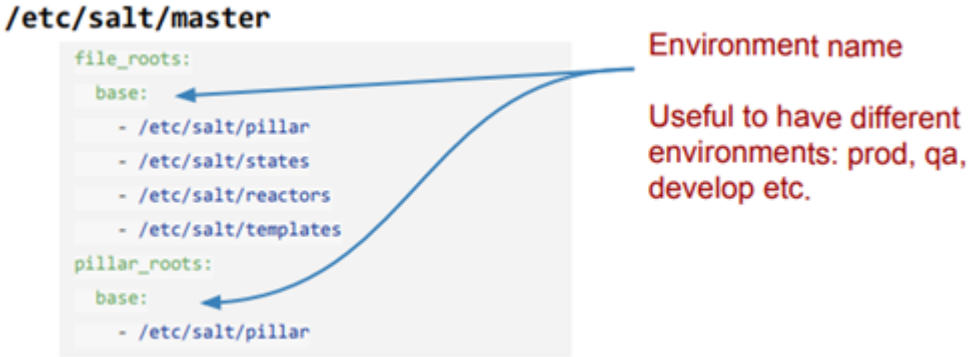

5. Save and quit by entering: wq
6. Restart the Salt-master file:

```
sudo systemctl start salt-master.service
sudo systemctl enable salt-master.service
```

9.2.3 Configuring the Salt-minion File

After the installation, modify the `/etc/salt/minion` configuration file as below:

1. Open the /etc/salt/minion file.
2. Replace #master: salt with master: 10.99.0.10.
3. Replace #hash_type: md5 with: hash_type: sha256.
4. Save and quit by entering: wq
5. Restart and enable Salt-minion.

```
sudo systemctl start salt-minion.service
```

9.2.4 Configuring the Proxy

1. Run /etc/salt/proxy.
2. Find the below attributes and fill them out as shown below:

/etc/salt/proxy

```
master: localhost
pki_dir: /etc/salt/pki/proxy
cachedir: /var/cache/salt/proxy
multiprocessing: False
mine_enabled: True
```

9.2.5 Creating the pillar Directory

1. Create a pillar directory under /etc/salt.

```
mkdir -r /etc/salt/pillar
```

2. Go to the /etc/salt/pillar directory
3. Create the top.sls file inside this directory.
   Per each switch, insert the following information:
   • DEVICE_ID
   • DEVICE_SLS_FILENAME
4. Create a new file: [DEVICE_SLS_FILENAME].sls
   Insert the following information into the above file:

```
proxy:
    proxytype: napalm
    driver: [DRIVER]
    host: [HOSTNAME]
    username: [USERNAME]
    passwd: [PASSWORD]
```

Example:

```
proxy:
```

```
    proxytype: napalm
    driver: onyx_ssh
    host: 10.209.37.247
    username: admin
    passwd: admin
  propt_name: switch20
  ssh_args:'-O PubkeyAuthentication=no'
```

5. Restart Salt on the server in order to use the new configuration

```
systemctl stop salt-minion
systemctl stop salt-master
systemctl stop salt-proxy@<switch_name>
systemctl start salt-master
systemctl start salt-minion
systemctl start salt-proxy@<switch_name>
```

9.2.6 Running Onyx Salt Commands on the Server

The following Salt command can be used:

1. Check if the switch is connected to the server running the Salt master:

```
salt onyx1 net.connected
```

2. Run any command on the switch using net.cli (example: using “show version”):

```
salt onyx1 net.cli 'show version'
```

3. Get the switch mac address:

```
salt onyx1 net.mac
```

4. Get the switch arp table:

```
salt onyx1 net.arp
```

5. Get switch information (uptime, vendor, os-version, etc):

```
salt onyx1 net.facts
```

6. Get the switch interfaces details:

```
salt onyx1 net.interfaces
```

9.3 Scheduled Jobs

The commands in this page may be used to manage and schedule the execution of jobs.

9.3.1 Commands

9.3.1.1 job

	job <job ID> no job <job ID> Creates a job. The no form of the command deletes the job.	
Syntax Description	job ID	Any integer
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # job 100` `switch (config job 100) #`	
Related Commands	show jobs	
Notes	• Job state is lost on reboot • If the command string includes a password, the configuration file will obscure the password and comment out the whole command	

9.3.1.2 command

	command <sequence #> | <command> no command <sequence #> Adds a CLI command to the job. The no form of the command deletes the command from the job.	
Syntax Description	sequence #	An integer that controls the order the command is executed relative to other commands in this job. The commands are executed in an ascending order.
	command	A CLI command
Default	N/A	
Configuration Mode	config job	
History	3.1.0000	
Example	`switch (config job 100) # command 10 “show images"`	
Related Commands	show jobs	
Notes	• The command must be defined with inverted commas (“”) • The command must be added as it was executed from the “config” mode. For example, in order to change the interface description you need to add the command: “interface <type> <number> description my-description”.	

9.3.1.3 comment

	comment <comment> no comment Adds a comment to the job. The no form of the command deletes the comment.	
Syntax Description	comment	A comment to be added to a specific job (string)
Default	N/A	
Configuration Mode	config job	
History	3.1.0000	
Example	`switch (config job 100) # comment Job_for_example`	
Related Commands	show jobs	
Notes		

9.3.1.4 enable

	enable no enable Enables the specified job. The no form of the command disables the specified job.
Syntax Description	N/A
Default	N/A
Configuration Mode	config job
History	3.1.0000
Example	`switch (config job 100) # enable`
Related Commands	show jobs
Notes	If a job is disabled, it will not be executed automatically according to its schedule; nor can it be executed manually.

9.3.1.5 execute

	execute Forces an immediate execution of the job.
Syntax Description	N/A
Default	N/A
Configuration Mode	config job
History	3.1.0000
Example	`switch (config job 100) # execute`
Related Commands	show jobs
Notes	• The job timer (if set) is not canceled and the job state is not changed: i.e. the time of the next automatic execution is not affected • The job will not be run if not currently enabled

9.3.1.6 fail-continue

	fail-continue no fail-continue Continues the job execution regardless of any job failures. The no form of the command returns fail-continue to its default.
Syntax Description	N/A
Default	A job will halt execution as soon as any of its commands fails
Configuration Mode	config job
History	3.1.0000
Example	`switch (config job 100) # fail-continue`
Related Commands	show jobs
Notes	

9.3.1.7 name

	name <job name> no name Configures a name for this job. The no form of the command resets the name to its default.	
Syntax Description	name	Specifies a name for the job (string)
Default	N/A	
Configuration Mode	config job	
History	3.1.0000	
Example	`switch (config job 100) # name my-job`	
Related Commands	show jobs	
Notes		

9.3.1.8 schedule type

	schedule type <recurrence type> no schedule type Sets the type of schedule the job will automatically execute on. The no form of the command resets the schedule type to its default.	
Syntax Description	recurrence type	The available schedule types are: • daily - the job is executed every day at a specified time • weekly - the job is executed on a weekly basis • monthly - the job is executed every month on a specified day of the month • once - the job is executed once at a single specified date and time • periodic - the job is executed on a specified fixed time interval, starting from a fixed point in time.
Default	once	
Configuration Mode	config job	

History	3.1.0000
Example	`switch (config job 100) # schedule type once`
Related Commands	show jobs
Notes	A schedule type is essentially a structure for specifying one or more future dates and times for a job to execute.

9.3.1.9 schedule <recurrence type>

	schedule <recurrence type> <interval and date> no schedule Sets the type of schedule the job will automatically execute on. The no form of the command resets the schedule type to its default.	
Syntax Description	recurrence type	The available schedule types are: • daily – the job is executed every day at a specified time • weekly – the job is executed on a weekly basis • monthly – the job is executed every month on a specified day of the month • once – the job is executed once at a single specified date and time • periodic – the job is executed on a specified fixed time interval, starting from a fixed point in time.
	interval and date	Interval and date, per recurrence type.
Default	once	
Configuration Mode	config job	
History	3.1.0000	
Example	`switch (config job 100) # schedule monthly interval 10`	
Related Commands	show jobs	
Notes	A schedule type is essentially a structure for specifying one or more future dates and times for a job to execute.	

9.3.1.10 show jobs

	show jobs [<job-id>] Displays configuration and state (including results of last execution, if any exist) of existing jobs.	
Syntax Description	job-id	A job ID whose information to display
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	

Example	switch (config) # show jobs 10 Job 10: Status: inactive Enabled: yes Continue on failure: no Schedule Type: once Time and date: 1970/01/01 00:00:00 +0000 Last Exec Time: Thu 2012/04/05 13:11:42 +0000 Next Exec Time: N/A Commands: Command 10: show terminal Last Output: CLI current session settings: Terminal width: 158 columns Terminal length: 38 rows Terminal type: xterm-256color X display setting: (none)
Related Commands	
Notes	

10 User Management, Authentication, & Security

- User Management & Security
- Cryptographic (X.509, IPSec) and Encryption

10.1 User Management & Security

10.1.1 User Accounts

There are two general user account types: admin and monitor. As admin, the user is privileged to execute all the available operations. As monitor, the user can execute operations that display system configuration and status, or set terminal settings.

User Role	Default Password
admin	admin
monitor	monitor

10.1.2 Authentication, Authorization, and Accounting (AAA)

AAA is a term describing a framework for intelligently controlling access to computer resources, enforcing policies, auditing usage, and providing the information necessary to bill for services. These combined processes are considered important for effective network management and security. The AAA feature allows you to verify the identity of, grant access to, and track the actions of users managing the system. The Remote Access Dial-In User Service (RADIUS) or Terminal Access Controller Access Control device Plus (TACACS+) or Lightweight Directory Access Protocol (LDAP) protocols are supported by the NVIDIA Onyx switch.

- Authentication—authentication provides the initial method of identifying each individual user, typically by entering a valid username and password before access is granted. The AAA server compares a user's authentication credentials with the user credentials stored in a database. If the credentials match, the user is granted access to the network or devices. If the credentials do not match, authentication fails and network access is denied.

- Authorization—following the authentication, a user must gain authorization for performing certain tasks. After logging into a system, for instance, the user may try to issue commands. The authorization process determines whether the user has the authority to issue such commands. Simply put, authorization is the process of enforcing policies: determining what types or qualities of activities, resources, or services a user is permitted. Usually, authorization occurs within the context of authentication. Once you have authenticated a user, they may be authorized for different types of access or activity.
- Accounting—the last level is accounting, which measures the resources a user consumes during access. This includes the amount of system time or the amount of data a user has sent and/or received during a session. Accounting is carried out by logging of session statistics and usage information, and is used for authorization control, billing, trend analysis, resource utilization, and capacity planning activities.

Authentication, authorization, and accounting services are often provided by a dedicated AAA server, a program that performs these functions. Network access servers interface with AAA servers using the Remote Authentication Dial-In User Service (RADIUS) protocol.

10.1.3 User Re-authentication

Re-authentication prevents users from accessing resources or perform tasks for which they do not have authorization. If credential information (e.g., AAA server information like IP address, key, port number, and so forth) that has been previously used to authenticate a user is modified, that user gets immediately logged out and then asked to re-authenticate.

10.1.4 RADIUS

RADIUS (Remote Authentication Dial-In User Service), widely used in network environments, is a client/server protocol and software that enables remote access servers to communicate with a central server to authenticate dial-in users and authorize their access to the requested system or service. It is commonly used for embedded network devices such as routers, modem servers, switches and so on. RADIUS is currently the de-facto standard for remote authentication. It is prevalent in both new and legacy systems.

It is used for several reasons:

- RADIUS facilitates centralized user administration
- RADIUS consistently provides some level of protection against an active attacker

10.1.5 TACACS+

TACACS (Terminal Access Controller Access Control System), widely used in network environments, is a client/server protocol that enables remote access servers to communicate with a central server to authenticate dial-in users and authorize their access to the requested system or service. It is commonly used for providing NAS (Network Access Security). NAS ensures secure access from remotely connected users. TACACS implements the TACACS Client and provides the AAA (Authentication, Authorization, and Accounting) functionalities.

TACACS is used for several reasons:

- Facilitates centralized user administration
- Uses TCP for transport to ensure reliable delivery
- Supports inbound authentication, outbound authentication and change password request for the authentication service
- Provides some level of protection against an active attacker

10.1.6 LDAP

LDAP (Lightweight Directory Access Protocol) is an authentication protocol that allows a remote access server to forward a user's log-on password to an authentication server to determine whether access can be allowed to a given system. LDAP is based on a client/server model. The switch acts as a client to the LDAP server. A remote user (the remote administrator) interacts only with the switch, not with the back-end server and database.

LDAP authentication consists of the following components:

- A protocol with a frame format that utilizes TCP over IP
- A centralized server that stores all the user authorization information
- A client: in this case, the switch

Each entry in the LDAP server is referenced by its Distinguished Name (DN). The DN consists of the user-account name concatenated with the LDAP domain name. The following is an example DN where the the user-account name is John:

```
uid=John,ou=people,dc=domain,dc=com
```

LDAP supports user membership in groups. If remote user is a member of admin or monitor group, it will be logged with admin or monitor capabilities respectively.
Supported group names for mapping are as follows:

- admin
- monitor

Supported group types (objectClass) on LDAP server side are as follows:

- groupOfNames
- posixGroup

10.1.7 System Secure Mode

System secure mode is a state that configures the switch system to run secure algorithms in compliance with FIPS 140-2 requirements. In this mode, unsecure algorithms are disabled and unsecure feature configurations are disallowed.

In this mode the system supports Federal Information Processing Standards (FIPS) 140-2, Security Requirements for Cryptographic Modules, which is a NIST (National Institute of Standards and Technology) publication that specifies the requirement for system cypher functionality.

When this mode is activated, all the modules which are used by the system are verified to work in compliance with the secure mode.

> Note that if system fails to load in secure mode it is loaded in non-secure mode.

Prerequisites:

1. Disable SNMPv1 and v2.

```
switch (config) # no snmp-server enable communities
```

2. Only allow SNMPv3 users with sha and aes-128.

```
switch (config) # snmp-server user <username> v3 auth sha <password1> priv aes-128 <password2>
```

3. Only allow SNMPv3 traps with sha and aes-128.

```
switch (config) # snmp-server host <ip-address> informs version 3 user <username> auth sha <password1> priv aes-128
 <password2>
```

4. Only allow SSHv2.

```
switch (config) # ssh server min-version 2
```

5. Enable SSH server strict security mode.

```
switch (config) # ssh server security strict
```

6. Disable HTTP access.

```
switch (config) # no web http enable
```

7. Enable HTTPS strict cyphers.

```
switch (config) # web https ssl ciphers TLS1.2
```

8. Disable router BGP neighbor password configuration.

```
switch (config) # no router bgp <as-number> neighbor <ip-address> password
```

9. Disable router BGP peer group password configuration.

```
switch (config) # no router bgp <as-number> peer-group <peer-group-name> password
```

10. Disable BGP password configuration.

```
switch (config) # no neighbor <ip-address> password
```

11. Disable MD5 password hashing on for users.

```
switch (config) # username <username> password <password>
```

> If a necessary prerequisite is not fulfilled the system does not activate secure mode and issues an advisory message accordingly.

To activate secure mode, do the following:

```
switch (config) # system secure-mode enable

Warning! Configuration is about to be saved and the system will be reloaded.
Type 'YES' to confirm the change in secure mode: YES
```

To deactivate secure mode, do the following:

```
switch (config) # no system secure-mode enable

Warning! Configuration is about to be saved and the system will be reloaded.
Type 'YES' to confirm the change in secure mode: YES
```

To verify secure mode configuration and state, do the following:

```
switch (config)# show system secure-mode

Secure mode configured: yes
Secure mode enabled: yes
```

10.1.8 User Management and Security Commands

- 10.1.8.1 User Accounts
  - 10.1.8.1.1 username
  - 10.1.8.1.2 show usernames
  - 10.1.8.1.3 show users
  - 10.1.8.1.4 show whoami
  - 10.1.8.1.5 password
  - 10.1.8.1.6 show password hardening
- 10.1.8.2 AAA Methods
  - 10.1.8.2.1 aaa accounting
  - 10.1.8.2.2 aaa authentication login
  - 10.1.8.2.3 aaa authentication attempts fail-delay
  - 10.1.8.2.4 aaa authentication attempts track
  - 10.1.8.2.5 aaa authentication attempts lockout
  - 10.1.8.2.6 aaa authentication attempts class-override
  - 10.1.8.2.7 aaa authentication attempts reset
  - 10.1.8.2.8 clear aaa authentication attempts
  - 10.1.8.2.9 aaa authorization
  - 10.1.8.2.10 show aaa
  - 10.1.8.2.11 show aaa authentication attempts
- 10.1.8.3 RADIUS
  - 10.1.8.3.1 radius-server
  - 10.1.8.3.2 radius-server enable
  - 10.1.8.3.3 radius-server host
  - 10.1.8.3.4 show radius
- 10.1.8.4 TACACS+
  - 10.1.8.4.1 tacacs-server

        - 10.1.8.4.2 tacacs-server enable
        - 10.1.8.4.3 tacacs-server host
        - 10.1.8.4.4 show tacacs
    - 10.1.8.5 LDAP
        - 10.1.8.5.1 ldap enable
        - 10.1.8.5.2 ldap base-dn
        - 10.1.8.5.3 ldap bind-dn/bind-password
        - 10.1.8.5.4 ldap group-attribute/group-dn
        - 10.1.8.5.5 ldap nested-group-search
        - 10.1.8.5.6 ldap nested-group-depth
        - 10.1.8.5.7 ldap nested-group-count
        - 10.1.8.5.8 ldap host
        - 10.1.8.5.9 ldap hostname-check enable
        - 10.1.8.5.10 ldap login-attribute
        - 10.1.8.5.11 ldap port
        - 10.1.8.5.12 ldap referrals
        - 10.1.8.5.13 ldap scope
        - 10.1.8.5.14 ldap ssl
        - 10.1.8.5.15 ldap timeout
        - 10.1.8.5.16 ldap version
        - 10.1.8.5.17 show ldap
        - 10.1.8.5.18 show ldap crl
    - 10.1.8.6 System Secure Mode
        - 10.1.8.6.1 system secure-mode enable
        - 10.1.8.6.2 show system secure-mode

10.1.8.1 User Accounts

10.1.8.1.1 username

	username <username> [capability <cap> | disable [login | password] | disconnect | full-name <name> | nopassword | password [0 | 7] <password>] no username <username> [capability | disable [login | password] | full-name] Creates a user and sets its capabilities, password and name. The no form of the command deletes the user configuration.	
Syntax Description	username	Specifies a username and creates a user account. New users are created initially with admin privileges but is disabled. Allowed characters for the username: • a-z • A-Z • 0-9 • period (.), underscore (_), hyphen (-) Any single character or combination of characters from the above is allowed except for a period "." in a single form.

<table>
<tr><td></td><td><table>
<tr><td>capability <cap></td><td>Defines user capabilities.
admin—full administrative capabilities
monitor—read only capabilities, can not change the running configuration
unpriv—can only query the most basic information, and cannot take any actions or change any configuration
v_admin—basic administrator capabilities</td></tr>
<tr><td>disable [login | password]</td><td>Disable—disable this account Disable login—disable all logins to this account Disable password—disable login to this account using a local password</td></tr>
<tr><td>disconnect</td><td>Logs out the specified user from the system.</td></tr>
<tr><td>name</td><td>Full name of the user.</td></tr>
<tr><td>nopassword</td><td>The next login of the user will not require password.</td></tr>
<tr><td>0 | 7</td><td>0—specifies a login password in cleartext 7—specifies a login password in encrypted text</td></tr>
<tr><td>password</td><td>Specifies a password for the user in string form. If [0 | 7] was not specified then the password is in cleartext.</td></tr>
</table></td></tr>
<tr><td>Default</td><td>The following usernames are available by default:
admin
monitor</td></tr>
<tr><td>Configuration Mode</td><td>config</td></tr>
<tr><td>History</td><td><table>
<tr><td>3.1.0000</td><td></td></tr>
<tr><td>3.4.0000</td><td>Updated example</td></tr>
<tr><td>3.4.1100</td><td>Updated example</td></tr>
<tr><td>3.6.2002</td><td>Added “disconnect” parameter</td></tr>
<tr><td>3.8.1000</td><td>Added "username" syntax description (allowed characters)</td></tr>
<tr><td>3.8.2000</td><td>Removed xmladmin and xmluser usernames due to XML depreciation</td></tr>
<tr><td>3.9.0900</td><td>Added note</td></tr>
</table></td></tr>
<tr><td>Example</td><td><pre>switch (config) # username monitor full-name smith</pre></td></tr>
<tr><td>Related Commands</td><td>show usernames
show users</td></tr>
<tr><td>Notes</td><td>To enable a user account, just set a password on it (or use the command “username <user> nopassword” to enable it with no password required for login)
Removing a user account does not terminate any current sessions that user has open; it just prevents new sessions from being established
Encrypted password is useful for the command “show configuration”, since the cleartext password cannot be recovered after it is set
The command "username <user> password <password>" or "username <user> password 0 <password>" are not security and will leave clear text in user's terminal (log and command history will be treated as sensitive information without clear text password). They are recommended to be replaced as "username <user> password" or "username <user> password" commands.</td></tr>
</table>

10.1.8.1.2 show usernames

	show usernames Displays list of users and their capabilities.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 | 3.8.1000 | Updated example output 3.8.2000 | Updated example output
Example	
switch (config) # show usernames USERNAME FULL NAME CAPABILITY ACCOUNT STATUS USERID System Administrator admin Local password login disabled admin System Administrator admin No password required for login monitor System Monitor monitor Password set (SHA512) root Root User admin No password required for login	
Related Commands	username show users
Notes	

10.1.8.1.3 show users

	show users [history] Displays logged in users and related information such as idle time and what host they have connected from.	
Syntax Description	history	Displays current and historical sessions.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example		

```
switch (config) # show users
USERNAME   FULL NAME                LINE    HOST               IDLE

admin      System Administrator     pts/0   172.22.237.174     0d0h34m4s
admin      System Administrator     pts/1   172.30.0.127       1d3h30m49s
admin      System Administrator     pts/3   172.22.237.34      0d0h0m0s

switch (config) #s how users history
admin      pts/3 172.22.237.34    Wed Feb  1 11:56   still logged in
admin      pts/3 172.22.237.34    Wed Feb  1 11:42 - 11:46  (00:04)
wtmp       begins                 Wed Feb  1 11:38:10 2012
```

Related Commands	username show usernames
Notes	

10.1.8.1.4 show whoami

	show whoami Displays username and capabilities of user currently logged in.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000

Example	`switch (config) # show whoami`
	`Current user: admin`
	`Capabilities: admin`

Related Commands	username show usernames show users
Notes	

10.1.8.1.5 password

	password [age expiration <days> | age warning <days> | history < length > | length minimal <length> | length maximal < length > | username-password-match enable | complexity-class <char class> | hardening enable] Configures restrictions for new passwords.	
Syntax Description	age expiration <days>	Specifies validity period of any password configured. Range: 0-365 days (0=password will not expire) Default: 365 days
	age warning <days>	Specifies how many days before expiration a warning message should be printed while logging in. Range: 0-30 days (0 indicates that a warning message will not be printed) Default: 15 days
	history < length >	Specifies how many passwords are saved per user. New password will be compared to previous passwords and will not be allowed if it is the same as an old one. Range: 0-20 passwords Default: 5 passwords
	length minimal <length>	Specifies minimal length of allowed password. Range: 1-32 characters Default: 8 characters
	length maximal < length>	Specifies maximal length of allowed password. Range: 64-80 characters Default: 64 characters
	username-password-match enable	Restricts user from having password identical to its username. Default: enabled The no form of this command will allow this.

	complexity-class <char class>	Specifies what characters must be used while configuring password. 1. none—no restrictions 2. lower 3. lower-upper 4. lower-upper-digit 5. lower-upper-digit-special Special characters allowed are: `~!@#$%^&*()-_=+[{}];:',<.> Default: lower-upper-digit
	hardening enable	Enable password restrictions. If enabled, all the above will be checked upon every new password that is being configured. Password that does not meet the requirements will be rejected. The no form will disable any password restrictions and every password will be allowed.
Default	Enabled. After upgrade, the feature will be disabled by default.	
Configuration Mode	Config	
History	3.9.2000	
Example	`switch (config) # password hardening enable`	
Related Commands	show password hardening	
Notes		

10.1.8.1.6 show password hardening

	show password hardening Displays all the configured password restrictions settings.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.9.2000

Example	`switch (config) # show password hardening`
	
	`Password settings:`
	`  Password hardening  : enabled`
	`  Min password length  : 8 (characters)`
	`  Max password length  : 64 (characters)`
	`  Character class  : Lowercase, uppercase and digits`
	`  Password history length  : 5`
	`  Different username and password: yes`
	`  Password aging  : enabled`
	`  Expiration warning message  : 15 (days)`
	`  Password age  : 365 (days)`
	
	`switch (config) # show password hardening`
	`Password settings:`
	`  Password hardening  : disabled`

Related Commands	password
Notes	• Wizard will prompt for enabling/disabling password hardening • Configuring password 7 while password hardening is enabled, will disable it

10.1.8.2 AAA Methods

10.1.8.2.1 aaa accounting

	aaa accounting changes default stop-only tacacs+ no aaa accounting changes default stop-only tacacs+ Enables logging of system changes to an AAA accounting server. The no form of the command disables the accounting.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Example	`switch (config) # aaa accounting changes default stop-only tacacs+`
Related Commands	show aaa
Notes	• TACACS+ is presently the only accounting service method supported • Change accounting covers both configuration changes and system actions that are visible under audit logging, however this feature operates independently of audit logging, so it is unaffected by the commands “logging level audit mgmt” or “configuration audit” • Configured TACACS+ servers are contacted in the order in which they appear in the configuration until one accepts the accounting data, or the server list is exhausted • Despite the name of the “stop-only” keyword, which indicates that this feature logs a TACACS+ accounting “stop” message, and in contrast to configuration change accounting, which happens after configuration database changes, system actions are logged when the action is started, not when the action has completed

10.1.8.2.2 aaa authentication login

	aaa authentication login default <auth method> [<auth method> [<auth method> [<auth method> [<auth method>]]]] no aaa authentication login Sets a sequence of authentication methods. Up to four methods can be configured. The no form of the command resets the configuration to its default.	
Syntax Description	auth-method	• local • radius • tacacs+ • ldap
Default	local	
Configuration Mode	Any command mode	
History	3.1.0000 3.7.1102—Updated notes	
Example	`switch (config) # aaa authentication login default  radius tacacs+ ldap local`	
Related Commands	show aaa	

Notes	• The order in which the methods are specified is the order in which the authentication is attempted. It is recommended that “local” is one of the methods selected. • When defining a remote server that to authenticate users against, once a connection is established with it, it does not go through other authentication methods. Meaning, if local is defined first, it will not go to other methods. If a remote server is defined first and then local (radius → local), then if the radius server is reachable, the response from this server will dictate whether the switch can be accessed or not (regardless of whether the user exists on any other authentication mothod).

10.1.8.2.3 aaa authentication attempts fail-delay

	aaa authentication attempts fail-delay <time> no aaa authentication attempts fail-delay Configures delay for a specific period of time after every authentication failure. The no form of the command resets the fail-delay to its default value.	
Syntax Description	time	Range: 0-60 seconds
Default	0	
Configuration Mode	config	
History	3.5.0200	
Example	`switch (config) # aaa authentication attempts fail-delay 1`	
Related Commands		
Notes		

10.1.8.2.4 aaa authentication attempts track

	aaa authentication attempts track {downcase | enable} no aaa authentication attempts track {downcase | enable} Configure tracking for failed authentication attempts. The no form of the command clears configuration for tracking authentication failures.	
Syntax Description	downcase	Does not convert all usernames to lowercase (for authentication failure tracking purposes only).
	enable	Disables tracking of failed authentication attempts.
Default	N/A	
Configuration Mode	config	
History	3.5.0200	
Example	`switch (config) # aaa authentication attempts track enable`	
Related Commands		
Notes	• This is required for the lockout functionality described below, but can also be used on its own for informational purposes. • Disabling tracking does not clear any records of past authentication failures, or the locks in the database. However, it does prevent any updates to this database from being made: no new failures are recorded. It also disables lockout, preventing new lockouts from being recorded and existing lockouts from being enforced.	

10.1.8.2.5 aaa authentication attempts lockout

	aaa authentication attempts lockout {enable | lock-time | max-fail | unlock-time} no aaa authentication attempts lockout {enable | lock-time | max-fail | unlock-time} Configures lockout of accounts based on failed authentication attempts. The no form of the command clears configuration for lockout of accounts based on failed authentication attempts.	
Syntax Description	enable	Enables locking out of user accounts based on authentication failures. This both suspends enforcement of any existing lockouts, and prevents any new lockouts from being recorded. If lockouts are later re-enabled, any lockouts that had been recorded previously resume being enforced; but accounts which have passed the max-fail limit in the meantime are NOT automatically locked at this time. They would be permitted one more attempt, and then locked, because of how the locking is done: lockouts are applied after an authentication failure, if the user has surpassed the threshold at that time. Lockouts only work if tracking is enabled. Enabling lockouts automatically enables tracking. Disabling tracking automatically disables lockouts.
	lock-time	Sets maximum permitted consecutive authentication failures before locking out users. Unlike the “max-fail” setting, this does take effect immediately for all accounts. If both unlock-time and lock-time are set, the unlock-time must be greater than the lock-time. This is not based on the number of consecutive failures, and is therefore divorced from most of the rest of the tally feature, except for the tracking of the last login failure.
	max-fail	Sets maximum permitted consecutive authentication failures before locking out users. This setting only impacts what lockouts are imposed while the setting is active; it is not retroactive to previous logins. So if max-fail is disabled or changed, this does not immediately cause any users to be changed from locked to unlocked or vice versa.
	unlock-time	Enables the auto-unlock of an account after a specified number of seconds if a user account is locked due to authentication failures, counting from the last valid login attempt. Unlike the “max-fail” setting, this does take effect immediately for all accounts. If both unlock-time and lock-time are set, the unlock-time must be greater than the lock-time. Careful with disabling the unlock-time, particularly if you have max-fail set to something, and have not overridden the behavior for the admin (i.e. they are subject to lockouts also). If the admin account gets locked out, and there are no other administrators who can aid, the user may be forced to boot single-user and use the pam_tallybyname command-line utility to unlock your account manually. Even if one is careful not to incur this many authentication failures, it makes the system more subject to DOS attacks.
Default	N/A	
Configuration Mode	config	

History	3.2.3000
Example	`switch (config) # aaa authentication attempts lockout enable`
Related Commands	
Notes	

10.1.8.2.6 aaa authentication attempts class-override

	aaa authentication attempts class-override {admin [no-lockout] | unknown {no-track | hash-username}} no aaa authentication attempts class-override {admin | unknown {no-track | hash-username}} Overrides the global settings for tracking and lockouts for a type of account. The no form of the command removes this override and lets the admin be handled according to the global settings.	
Syntax Description	admin	Overrides the global settings for tracking and lockouts for the admin account. This applies only to the single account with the username “admin”. It does not apply to any other users with administrative privileges.
	no-lockout	Prevents the admin user from being locked out though authentication failure history is still tracked (if tracking is enabled overall).
	unknown	Overrides the global settings for tracking and lockouts for unknown accounts. The “unknown” class here contains the following categories: • Real remote usernames which simply failed authentication • Mis-typed remote usernames • Passwords accidentally entered as usernames • Bogus usernames made up as part of an attack on the system
	hash-username	Applies a hash function to the username and stores the hashed result in lieu of the original
	no-track	Does not track authentication for such users (which of course also implies no-lockout)
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Example	`switch (config) # aaa authentication attempts class-override admin no-lockout`	
Related Commands		
Notes		

10.1.8.2.7 aaa authentication attempts reset

	aaa authentication attempts reset {all | user <username>} [{no-clear-history | no-unlock}] Clears the authentication history for and/or unlocks specified users.	
Syntax Description	all	Applies function to all users
	user	Applies function to a specific user

	no-clear-history	Leaves the history of login failures but unlocks the account
	no-unlock	Leaves the account locked but clears the history of login failures
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Example	`switch (config) # aaa authentication attempts reset user admin all`	
Related Commands		
Notes		

10.1.8.2.8 clear aaa authentication attempts

	clear aaa authentication attempts {all | user <username>} [no-clear-history | no-unlock] Clears the authentication history for and/or unlocks specified users.	
Syntax Description	all	Applies function to all users.
	user	Applies function to a specific user.
	no-clear-history	Clears the history of login failures.
	no-unlock	Unlocks the account.
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Example	`switch (config) # aaa authentication attempts reset user admin no-clear-history`	
Related Commands		
Notes		

10.1.8.2.9 aaa authorization

	aaa authorization map [default-user <username> | order <policy> | fallback] no aaa authorization map [default-user | order | fallback] Sets the mapping permissions of a user in case a remote authentication is done. The no form of the command resets the attributes to default.	
Syntax Description	username	Specifies what local account the authenticated user will be logged on as when a user is authenticated (via RADIUS or TACACS+ or LDAP) and does not have a local account. If the username is local, this mapping is ignored.

	order <policy>	Sets the user mapping behavior when authenticating users via RADIUS or TACACS+ or LDAP to one of three choices. The order determines how the remote user mapping behaves. If the authenticated username is valid locally, no mapping is performed. The setting has the following three possible behaviors: • local-only—maps all remote users to the user specified by the command “aaa authorization map default-user <user name>”. Any vendor attributes received by an authentication server are ignored. • remote-first—if a local-user mapping attribute is returned and it is a valid local username, it maps the authenticated user to the local user specified in the attribute. Otherwise, it uses the user specified by the default-user command. • remote-only—maps a remote authenticated user if the authentication server sends a local-user mapping attribute. If the attribute does not specify a valid local user, no further mapping is tried.
	fallback	Sets the authenticating fallback behavior via RADIUS or TACACS+ or LDAP. This option attempts to authenticate username through the next authentication method listed in case of an error. • server-err—performs fallback if an error occurs while connecting to remote AAA server (e.g., server is down, not responding, and so forth)
Default	Default user—admin Map order—remote-first Order fallback—server-err	
Configuration Mode	config	
History	3.1.0000 3.7.1000—Added “fallback” parameter 3.7.1000—Updated syntax	
Example	`switch (config) # aaa authorization map default-user admin`	
Related Commands	show aaa username	
Notes	• If, for example, the user is locally defined to have admin permission, but in a remote server such as RADIUS the user is authenticated as monitor and the order is remote-first, then the user is given monitor permissions. • The user must be careful when disabling AAA authorization map fallback server-err, because if the remote server stops working then the user may lock themselves out. • If AAA authorization order policy is configured to remote-only, then when upgrading to 3.4.3000 or later from an older version, this policy is changed to remote-first.	

10.1.8.2.10 show aaa

	show aaa Displays the AAA configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.7.0020—Example updated

Example	switch (config) # show aaa AAA authorization: Default User: admin Map Order: remote-first Fallback on server-err: yes Authentication method(s): local Accounting method(s): tacacs+
Related Commands	aaa accounting aaa authentication aaa authorization show aaa show usernames username
Notes	

10.1.8.2.11 show aaa authentication attempts

	show aaa authentication attempts [configured | status user <username>]] Displays the current authentication, authorization and accounting settings.	
Syntax Description	authentication attempts	Displays configuration and history of authentication failures.
	configured	Displays configuration of authentication failure tracking.
	status user	Displays status of authentication failure tracking and lockouts for specific user.
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.1000 3.5.0200—Updated example	

Example

```
switch (config) # show aaa authentication attempts

Configuration for authentication failure tracking and locking:

   Track authentication failures:                   yes
   Lock accounts based on authentication failures:  yes
   Override treatment of 'admin' user:              (none)
   Override treatment of unknown usernames:         hash-usernames
   Convert usernames to lowercase for tracking:     no
   Delay after each auth failure (fail delay):      none



Configuration for lockouts based on authentication failures:
   Lock account after consecutive auth failures:    5
   Allow retry on locked accounts (unlock time):    after 15 second(s)
   Temp lock after each auth failure (lock time):   none

Username                                  Known  Locked  Failures  Last fail time         Last fail from
--------                                  -----  ------  --------  --------------         --------------
0Q72B43EHBKT8CB5AF5PGRX3U3B3TUL4CYJP93N(*)  no     no      1         2020/05/20 14:29:19    ttyS0

(*) Hashed for security reasons
```

Related Commands	
Notes	

10.1.8.3 RADIUS

10.1.8.3.1 radius-server

	radius-server {key <secret>| retransmit <retries> | timeout <seconds>} no radius-server {key | retransmit | timeout} Sets global RADIUS server attributes. The no form of the command resets the attributes to their default values.	
Syntax Description	secret	Sets a secret key (shared hidden text string), known to the system and to the RADIUS server.
	retries	Number of retries (0-5) before exhausting from the authentication.
	seconds	Timeout in seconds between each retry (1-60).
Default	3 seconds, 1 retry	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # radius-server retransmit 3`	
Related Commands	aaa authorization radius-server host show radius	
Notes	Each RADIUS server can override those global parameters using the command “radius-server host”.	

10.1.8.3.2 radius-server enable

	radius-server [vrf <vrf-name>] enable [force] no radius-server [vrf <vrf-name>] enable Enables RADIUS in VRF. The no form of the command disables RADIUS in a specified VRF.	
Syntax Description	vrf-name	VRF name
	force	Enables RADIUS in the specified VRF and sets all relevant RADIUS option to default
Default	RADIUS is enabled by default	
Configuration Mode	config	
History	3.9.2000	
Example		
Related Commands	`switch (config) # radius-server vrf mgmt enable`	
Notes	If VRF management exists, RADIUS will be enabled on VRF management. If VRF management not does not exist, RADIUS will be enabled on VRF default.	

10.1.8.3.3 radius-server host

	radius-server host <IP address> [enable | auth-port <port> | key <secret> | prompt-key | retransmit <retries> | timeout <seconds>| cipher <none | eap-peap>] no radius-server host <IP address> [auth-port | enable | cipher] Configures RADIUS server attributes. The no form of the command resets the attributes to their default values and deletes the RADIUS server.	
Syntax Description	IP address	RADIUS server IP address
	enable	Administrative enable of the RADIUS server
	auth-port	Configures authentication port to use with this RADIUS server
	port	RADIUS server UDP port number
	key	Configures shared secret to use with this RADIUS server
	prompt-key	Prompt for key, rather than entering on command line
	retransmit	Configures retransmit count to use with this RADIUS server
	retries	Number of retries (0-5) before exhausting from the authentication
	timeout	Configures timeout between each try
	seconds	Timeout in seconds between each retry (1-60)
	cipher	Configures which cipher to use for communication encryption <none | eap-peap>
Default	3 seconds, 1 retry Default UDP port is 1812	
Configuration Mode	config	
History	3.1.0000 3.8.1000—Updated command description, syntax description & example	
Example	`switch (config) # radius-server host fe80::202:b3ff:fe1e:8329` `switch (config) # radius-server host 40.40.40.40`	
Related Commands	aaa authorization radius-server show radius	
Notes	• RADIUS servers are tried in the order they are configured • If you do not specify a parameter for this configured RADIUS server, the configuration will be taken from the global RADIUS server configuration. Refer to the command “radius-server”.	

10.1.8.3.4 show radius

	show radius Displays RADIUS configurations.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode

History	3.1.0000 3.6.6000—Updated example 3.8.1000—Updated command description, syntax description & example 3.9.2000—Updated example , adding the "administratively" and "VRF name" fields

Example	`switch (config) # show radius`
	
	`RADIUS defaults:`
	`  administratively: enabled`
	`  VRF name:  : mgmt`
	`  Key  : ********`
	`  Timeout  : 3`
	`  Retransmit  : 1`
	
	`RADIUS servers:`
	` 1.1.1.1:1812 :`
	` Enabled  : yes`
	` Key  : ********`
	` Timeout  : 3 (default)`
	` Retransmit  : 1 (default)`
	` Cipher  : none`
	
	`40.40.40.40:1812:`
	` Enabled  : yes`
	` Key  : ********`
	` Timeout  : 3 (default)`
	` Retransmit  : 1 (default)`

Related Commands	aaa authorization radius-server radius-server host
Notes	

10.1.8.4 TACACS+

10.1.8.4.1 tacacs-server

	tacacs-server {key <secret>| retransmit <retries> | timeout <seconds>} no tacacs-server {key | retransmit | timeout} Sets global TACACS+ server attributes. The no form of the command resets the attributes to default values.	
Syntax Description	secret	Set a secret key (shared hidden text string), known to the system and to the TACACS+ server.
	retries	Number of retries (0-5) before exhausting from the authentication.
	seconds	Timeout in seconds between each retry. Reang: 1-60
Default	3 seconds, 1 retry	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # tacacs-server retransmit 3`	
Related Commands	aaa authorization show radius show tacacs tacacs-server host	
Notes	Each TACACS+ server can override those global parameters using the command “tacacs-server host”.	

10.1.8.4.2 tacacs-server enable

	tacacs-server [vrf <vrf-name>] enable [force] no tacacs-server [vrf <vrf-name>] enable Enables TACACS in VRF. The no form of the command disables TACACS in a specified VRF.	
Syntax Description	vrf-name	VRF name
	force	Enables TACACS in the specified VRF and sets all relevant TACACS option to default
Default	TACACS is enabled by default	
Configuration Mode	config	
History	3.9.2000	
Example	`switch (config) # tacacs-server vrf mgmt enable`	
Related Commands		
Notes	If VRF management exists, TACACS will be enabled on VRF management. If VRF management not does not exist, TACACS will be enabled on VRF default.	

10.1.8.4.3 tacacs-server host

	tacacs-server host <IP address> {enable | auth-port <port> | auth-type <type> | key <secret> | prompt-key | retransmit <retries> | timeout <seconds>} no tacacs-server host <IP address> {enable | auth-port} Configures TACACS+ server attributes. The no form of the command resets the attributes to their default values and deletes the TACACS+ server.	
Syntax Description	IP address	TACACS+ server IP address.
	enable	Administrative enable for the TACACS+ server.
	auth-port	Configures authentication port to use with this TACACS+ server.
	port	TACACS+ server UDP port number.
	auth-type	Configures authentication type to use with this TACACS+ server.
	type	Authentication type. Possible values are: • ASCII • PAP (Password Authentication Protocol)
	key	Configures shared secret to use with this TACACS+ server.
	secret	Sets a secret key (shared hidden text string), known to the system and to the TACACS+ server.
	prompt-key	Prompts for key, rather than entering key on command line.
	retransmit	Configures retransmit count to use with this TACACS+ server.
	retries	Number of retries (0-5) before exhausting from the authentication.
	timeout	Configures timeout to use with this TACACS+ server.
	seconds	Timeout in seconds between each retry. Range: 1-60

Default	3 seconds, 1 retry Default TCP port is 49 Default auth-type is PAP
Configuration Mode	config
History	3.1.0000
Example	`switch (config) # tacacs-server host 40.40.40.40`
Related Commands	aaa authorization show tacacs tacacs-server
Notes	• TACACS+ servers are tried in the order they are configured • A PAP auth-type similar to an ASCII login, except that the username and password arrive at the network access server in a PAP protocol packet instead of being typed in by the user, so the user is not prompted • If the user does not specify a parameter for this configured TACACS+ server, the configuration will be taken from the global TACACS+ server configuration. Refer to the command “tacacs-server”.

10.1.8.4.4 show tacacs

	show tacacs Displays TACACS+ configurations.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.6.6000—Updated example 3.9.2000—Updated example , adding the "administratively" and "VRF name" fields

Example	`switch (config) # show tacacs`
	
	`TACACS+ defaults:`
	`  Key : ********`
	`  Timeout : 3 Retransmit: 1`
	`switch (config) # show tacacs`
	
	`TACACS+ defaults:`
	`  administratively: enabled`
	`  VRF name:  : mgmt`
	`  Key  :  ********`
	`  Timeout  :  3`
	`  Retransmit  :  1`
	`TACACS+ servers:`
	`  1.1.1.1:49:`
	`  Enabled  : yes`
	`  Auth Type : pap`
	`  Key  : ********`
	`  Timeout  : 3 (default)`
	`  Retransmit: 1 (default)`

Related Commands	aaa authorization tacacs-server tacacs-server host
Notes	

10.1.8.5 LDAP

10.1.8.5.1 ldap enable

	ldap [vrf <vrf-name>] enable [force] no ldap [vrf <vrf-name>] enable Enables LDAP in VRF. The no form of the command disables LDAP in a specified VRF.	
Syntax Description	force	Enables LDAP in the specified VRF while setting all relevant LDAP options to default.
Default	LDAP enabled	
Configuration Mode	config	
History	3.9.2000	
Example	`switch (config) # ldap vrf mgmt enable`	
Related Commands		
Notes	If VRF mgmt exists, LDAP will be enabled on VRF mgmt. If there is no VRF mgmt, LDAP will be enabled on the "default" VRF.	

10.1.8.5.2 ldap base-dn

	ldap base-dn <string> no ldap base-dn Sets the base distinguished name (location) of the user information in the schema of the LDAP server. The no form of the command resets the attribute to its default values.	
Syntax Description	string	A case-sensitive string that specifies the location in the LDAP hierarchy where the server should begin searching when it receives an authorization request. For example: “ou=users,dc=example,dc=com”, with no spaces. Where: • ou—Organizational unit • dc—Domain component • cn—Common name • sn—Surname
Default	ou=users,dc=example,dc=com	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated example
Example	`switch (config) # ldap base-dn ou=department,dc=example,dc=com`	
Related Commands	show ldap	
Notes		

10.1.8.5.3 ldap bind-dn/bind-password

	ldap {bind-dn | bind-password} <string> no ldap {bind-dn | bind-password} Gives the distinguished name or password to bind to on the LDAP server. This can be left empty for anonymous login (the default). The no form of the command resets the attribute to its default values.	
Syntax Description	string	A case-sensitive string that specifies distinguished name or password to bind to on the LDAP server.
Default	“”	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated example
Example	`switch (config) # ldap bind-dn my-dn` `switch (config) # ldap bind-password my-password`	
Related Commands	show ldap	
Notes	For anonymous login, bind-dn and bind-password should be empty strings “”.	

10.1.8.5.4 ldap group-attribute/group-dn

	ldap {group-attribute {<group-att> |member | uniqueMember} | group-dn <group-dn>} no ldap {group-attribute | group-dn} Sets the distinguished name or attribute name of a group on the LDAP server. The no form of the command resets the attribute to its default values.	
Syntax Description	group-att	Specifies a custom attribute name.
	member	groupOfNames or group membership attribute.
	uniqueMember	groupOfUniqueNames membership attribute.
	group-dn	DN of group required for authorization.
Default	group-att: member group-dn: “”	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated example
Example	`switch (config) # ldap group-attribute member` `switch (config) # ldap group-dn my-group-dn`	
Related Commands	show ldap	
Notes	• The user’s distinguished name must be listed as one of the values of this attribute, or the user will not be authorized to log in • After login authentication, if the group-dn is set, a user must be a member of this group or the user will not be authorized to log in. If the group is not set (“”—the default) no authorization checks are done.	

10.1.8.5.5 ldap nested-group-search

	ldap nested-group-search no ldap nested-group-search Enable LDAP nested-group search mechanism for user-authentication group matching. The no form of the command resets the attribute to its default values.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.10.2000
Example	`switch (config) # ldap nested-group-search` `switch (config) # no ldap nested-group-search`
Related Commands	ldap nested-group-depth ldap nested-group-count show ldap
Notes	

10.1.8.5.6 ldap nested-group-depth

	ldap nested-group-depth <1-9> no ldap nested-group-depth Sets LDAP maximum depth for nested-group search. The no form of the command resets search depth to default (3).
Syntax Description	N/A
Default	3
Configuration Mode	config
History	3.10.2000
Example	`switch (config) # ldap nested-group-depth 6` `switch (config) # no ldap nested-group-depth`
Related Commands	ldap nested-group-search ldap nested-group-count show ldap
Notes	

10.1.8.5.7 ldap nested-group-count

	ldap nested-group-count <1-10000> no ldap nested-group-count Sets LDAP maximum number of queried nested-groups. The no form of the command resets search depth to default (1000).
Syntax Description	N/A
Default	1000
Configuration Mode	config
History	3.10.2000

Example	`switch (config) # ldap nested-group-count 500` `switch (config) # no ldap nested-group-count`
Related Commands	ldap nested-group-depth ldap nested-group-search show ldap
Notes	

10.1.8.5.8 ldap host

	ldap host <ip-address> [order <number> last] no ldap host <ip-address> Adds an LDAP server to the set of servers used for authentication. The no form of the command deletes the LDAP host.	
Syntax Description	ip-address	IPv4 or IPv6 address.
	number	The order of the LDAP server.
	last	The LDAP server will be added in the last location.
Default	No hosts configured	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated example
Example	`switch (config) # ldap host 10.10.10.10`	
Related Commands	show aaa show ldap	
Notes	• The system will select the LDAP host to try according to its order • New servers are by default added at the end of the list of servers	

10.1.8.5.9 ldap hostname-check enable

	ldap hostname-check enable no ldap hostname-check enable Enables LDAP hostname check. The no form of the command disables LDAP hostname check.
Syntax Description	N/A
Default	No hosts configured
Configuration Mode	config
History	3.6.8008
Example	`switch (config) # ldap hostname-check enable`
Related Commands	show aaa show ldap
Notes	

10.1.8.5.10 ldap login-attribute

	ldap login-attribute {<string> | uid | sAMAccountName} no ldap login-attribute Sets the attribute name which contains the login name of the user. The no form of the command resets this attribute to its default.	
Syntax Description	string	Custom attribute name.
	uid	LDAP login name is taken from the user login username.
	sAMAccountName	SAM Account name, active directory login name.
Default	sAMAccountName	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated example
Example	`switch (config) # ldap login-attribute uid`	
Related Commands	show aaa show ldap	
Notes		

10.1.8.5.11 ldap port

	ldap port <port> no ldap port Sets the TCP port on the LDAP server to connect to for authentication. The no form of the command resets this attribute to its default value.	
Syntax Description	port	TCP port number
Default	389	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated example
Example	`switch (config) # ldap port 1111`	
Related Commands	show aaa show ldap	
Notes		

10.1.8.5.12 ldap referrals

	ldap referrals no ldap referrals Enables LDAP referrals. The no form of the command disables LDAP referrals.
Syntax Description	N/A
Default	LDAP referrals are enabled

Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated example
Example	`switch (config) # no ldap referrals`	
Related Commands	show aaa show ldap	
Notes	Referral is the process by which an LDAP server, instead of returning a result, will return a referral (a reference) to another LDAP server which may contain further information.	

10.1.8.5.13 ldap scope

	ldap scope <scope> no ldap scope Specifies the extent of the search in the LDAP hierarchy that the server should make when it receives an authorization request. The no form of the command resets the attribute to its default value.	
Syntax Description	scope	• one-level—searches the immediate children of the base dn • subtree—searches at the base DN and all its children
Default	subtree	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated example
Example	`switch (config) # ldap scope subtree`	
Related Commands	show aaa show ldap	
Notes		

10.1.8.5.14 ldap ssl

	ldap ssl {ca-list <options> | cert-verify | ciphers {all | TLS1.2} | crl-check {enable | file fetch all [vrf <vrf-name>] <path>} | mode <mode> | port <port-number>} no ldap ssl {cert-verify | ciphers | crl-check enable | mode | port} Sets SSL parameter for LDAP. The no form of the command resets the attribute to its default value.

Syntax Description	options	This command specifies the list of supplemental certificates of authority (CAs) from the certificate configuration database that is to be used by LDAP for authentication of servers when in TLS or SSL mode. The options are: • default-ca-list—uses default supplemental CA certificate list • none—no supplemental list, uses the built-in one only CA certificates are ignored if “ldap ssl mode” is not configured as either “tls” or “ssl”, or if “no ldap ssl cert-verify” is configured. The default-ca-list is empty in the factory default configuration. Use the command: “crypto certificate ca-list default-ca-list name” to add trusted certificates to that list. The “default-ca-list” option requires LDAP to consult the system’s configured global default CA-list for supplemental certificates.
	cert-verify	Enables verification of SSL/TLS server certificates. This may be required if the server's certificate is self-signed, or does not match the name of the server.
	ciphers {all | TLS1.2}	Sets SSL mode to be used
	crl-check enable	Enables LDAP CRL check
	crl-check file fetch	Fetches CRL from remote server. CRL must be a valid PEM file unless a proper message shown. Supported formats: SCP, HTTP, HTTPS, FTP, and FTPS.
	mode	Sets the security mode for connections to the LDAP server. • none—requests no encryption for the LDAP connection • ssl—the SSL-port configuration is used, an SSL connection is made before LDAP requests are sent (LDAP over SSL) • start-tls—the normal LDAP port is used, an LDAP connection is initiated, and then TLS is started on this existing connection
	vrf-name	VRF to be affected. If "vrf-name" parameter is not specified, "default" VRF will be used.
	port-number	Sets the port on the LDAP server to connect to for authentication when the SSL security mode is enabled (LDAP over SSL)
Default	cert-verify—enabled mode—none (LDAP SSL is not activated) port-number—636 ciphers—all	
Configuration Mode	config	
History	3.1.0000	
	3.2.3000	Added ca-list argument
	3.4.0000	Added “ssl ciphers” parameter and Updated example
	3.6.8008	Added the parameter “crl-check”
	3.9.2000	Addded VRF option

Example	`switch (config) # ldap ssl crl-check file fetch scp://root:pass@1.1.1.1/etc/pki/crl.pem`
	
	`100.0% [##########################################################################]`	

Related Commands	show aaa show ldap	

Notes	• If available, the TLS mode is recommended, as it is standardized, and may also be of higher security • The port number is used only for SSL mode. If the security mode selected is TLS, the LDAP port number is used.

10.1.8.5.15 ldap timeout

	ldap {timeout-bind | timeout-search} <seconds> no ldap {timeout-bind | timeout-search} Sets a global communication timeout in seconds for all LDAP servers to specify the extent of the search in the LDAP hierarchy that the server should make when it receives an authorization request. The no form of the command resets the attribute to its default value.	
Syntax Description	timeout-bind	Sets the global LDAP bind timeout for all LDAP servers.
	timeout-search	Sets the global LDAP search timeout for all LDAP servers.
	seconds	Number of seconds. Range: 1-60
Default	5 seconds	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated example
Example	`switch (config) # ldap timeout-bind 10`	
Related Commands	show aaa show ldap	
Notes		

10.1.8.5.16 ldap version

	ldap version <version> no ldap version Sets the LDAP version. The no form of the command resets the attribute to its default value.	
Syntax Description	version	Sets the LDAP version Available values: 2, 3
Default	3	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Updated example
Example	`switch (config) # ldap version 3`	
Related Commands	show aaa show ldap	
Notes		

10.1.8.5.17 show ldap

	show ldap Displays LDAP configurations.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.4.0000	Updated example
	3.6.8008	Updated example
	3.10.2000	Updated example to reflect the following added fields: "Nested-group search," "nested-group search depth," and "nested-search maximum group count"

Example	`switch (config) # show ldap`
	
	`administratively : enabled`
	`VRF name: : mgmt`
	`User base DN : ou=users,dc=example,dc=com`
	`User search scope : subtree`
	`Login attribute : sAMAccountName`
	`Bind DN :`
	`Bind password : ********`
	`Group base DN :`
	`Group attribute : member`
	`Nested-group search : disabled`
	`Nested-group search depth : 3`
	`Nested-search maximum group count: 1000`
	`LDAP version : 3`
	`Referrals : yes`
	`Server port : 389`
	`Search Timeout : 5`
	`Bind Timeout : 5`
	`Server Hostname check : no`
	`SSL mode : none`
	`Server SSL port : 636 (not active)`
	`SSL ciphers : all (not active)`
	`SSL cert verify : yes`
	`SSL ca-list : default-ca-list`
	`SSL CRL check : no`
	
	`LDAP servers:`
	`No LDAP servers configured.`	

Related Commands	show aaa show ldap	
Notes		

10.1.8.5.18 show ldap crl

	show ldap crl Displays current CRL configured by the user.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.8008

Example	switch (config) # show ldap crl -----BEGIN CERTIFICATE----- MIIDVzCSd..... -----END CERTIFICATE-----
Related Commands	show aaa show ldap
Notes	

10.1.8.6 System Secure Mode

10.1.8.6.1 system secure-mode enable

	system secure-mode enable no system secure-mode enable Enables secure mode on the switch. The no form of the command disables secure mode.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.5.0200 3.10.2000: Added note
Example	switch (config) # system secure-mode enable Warning! Configuration is about to be saved and the system will be reloaded. Type 'YES' to confirm the change in secure mode: YES
Related Commands	user <username> password <password> ssh server min-version ssh server security strict snmp-server user no neighbor <ip-address> password ntp server disable ntp server keyID router bgp neighbor password router bgp peer-group password

Notes	• Before enabling secure mode, the command performs the following configuration checks:
	  • NTP Key ID cannot be MD5 when secure mode is enabled
	  • SSH min-version cannot be 1 when enabling secure mode
	  • SSH security must be set to strict security
	  • SNMPv3 user auth cannot be md5 when enabling secure mode
	  • SNMPv3 user priv cannot be des when enabling secure mode
	  • SNMPv3 trap auth cannot be md5 when enabling secure mode
	  • SNMPv3 trap priv cannot be des when enabling secure mode
	  • Router BGP neighbor password cannot be set when enabling secure mode
	  • Router BGP peer-group password cannot be set when enabling with secure mode
	  • User password hash cannot be MD5 when secure mode is enabled
	  Only if the check passes, secure mode is enabled on the switch system.
	• When secure mode is enabled extra reboot may happen after next steps: install new image and boot to newly installed image.

10.1.8.6.2 show system secure-mode

	show system secure-mode Displays the security mode of the switch system.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.4.2300

Example	`switch (config) # show system secure-mode`
	
	`Secure mode configured: yes`
	`Secure mode enabled : yes`

Related Commands	system secure-mode enable
Notes	• “Secure mode configuration” describes the user configuration • “Secure mode enabled” describes the system state

10.1.9 802.1x Protocol

The 802.1x (dot1x) standard describes a way to authenticate hosts (or supplicants) and to allow connection only to a list of allowed hosts pre-configured on an authentication server. The authentication is performed by the switch (authenticator) which negotiates the authentication with a RADIUS server (authentication server). This allows to block traffic from non-authenticated sources.

The 802.1x protocol defines the following roles:

- Supplicant – the host. It provides the authentication credentials to the authenticator and awaits approval.
- Authenticator – the device that connects the supplicant to the network, and checks the authentication with the authentication server. The authenticator is also in charge of blocking and isolating of new client till authenticated and allowing communication once the client has passed the authentication. The switch acts as an authenticator.
- Authentication server – a RADIUS server which can authenticate the user.

The 802.1x is available only on access physical ports. It is not available on LAG and MLAG ports.

A local analyzer port cannot support 802.1x protocol.

802.1x cannot be activated on router port interfaces.

802.1x cannot run on a port configured to switchport trunk or hybrid.

Management interfaces cannot be configured as 802.1x port access entity (PAE) authenticators.

10.1.9.1 802.1x Operating Modes

The following operating modes are supported in 802.1x:

- Single host – only one supplicant can communicate through the port.Once authentication of the supplicant is accepted by the authentication server, the switch allows it access. If the supplicant logs off or the port state is changed, the port becomes unauthenticated. And if a different supplicant tries to access through this port, its bidirectional traffic is discarded (including authentication traffic).

  An exception to this is multicast and broadcast traffic which do get transmitted over the interface once authenticated and are exposed to an unauthorized supplicant if it exists.

- Multi-host mode – allows connection of multiple hosts over a single port. Only the first supplicant is authenticated. Subsequent hosts have network access without the need to authenticate.

10.1.9.2 Configuring 802.1x

1. Enable 802.1x protocol.

```
switch (config) # protocol dot1x
```

2. Enable the system as authenticator.

```
switch (config) # dot1x system-auth-control
```

3. Configure RADIUS server parameters.

```
switch (config) # dot1x radius-server host 10.10.10.10 key my4uth3nt1c4t10nk3y retransmit 2 timeout 3
```

4. Enter the configuration mode of an Ethernet interface.

```
switch (config) # interface ethernet 1/1
switch (config interface ethernet 1/1) #
```

5. Configure the interface as a port access entity authenticator.

```
switch (config interface ethernet 1/1) # dot1x pae authenticator
```

6. Configure the interface to perform authentication on ingress traffic.

```
switch (config interface ethernet 1/1) # dot1x port-control auto
```

7. Verify 802.1x configuration.

```
switch (config interface ethernet 1/1) # show dot1x interfaces ethernet 1/1

Eth1/1
  PAE Status:                       Enabled
  Configured host mode:             Multi-host
  Configured port-control:          Auto
  Authentication status:            Unauthorized
  Re-Authentication:                Disabled
  Re-Authentication period (sec):   -
  Tx wait period (sec):             30
  Quiet period (sec):               60
  Max request retry:                2
  Last EAPOL RX source MAC:         00:00:00:00:00:00
```

10.1.9.3 Dot1x Commands

10.1.9.3.1 protocol dot1x

	protocol dot1x no protocol dot1x Enables 802.1x EAPOL protocol. The no form of the command disables 802.1x EAPOL protocol.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.4.2008
Example	`switch (config)# protocol dot1x`
Related Commands	
Notes	

10.1.9.3.2 dot1x clear-statistics

	dot1x clear-statistics Resets the 802.1x counters on all or a specific port.
Syntax Description	N/A
Default	N/A
Configuration Mode	config config interface ethernet
History	3.4.2008
Example	`switch (config)# dot1x clear-statistics`
Related Commands	
Notes	

10.1.9.3.3 dot1x pae authenticator

	dot1x pae authenticator no dot1x pae authenticator Configures the port as a 802.1x port access entity (PAE) authenticator. The no form of the command disables the port from being a 802.1x PAE authenticator.

Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface ethernet
History	3.4.2008
Example	`switch (config interface ethernet 1/2)# dot1x system-auth-control`
Related Commands	
Notes	

10.1.9.3.4 dot1x host-mode

	dot1x host-mode [multi-host | single-host] no dot1x host-mode Configures the authentication mode to either multi-host or single-host. The no form of the command resets the parameter to its default.	
Syntax Description	multi-host	Sets the interface to operate in a port-based mode
	single-host	Sets the interface to operate in a MAC-based mode with support of a single supplicant per interface
Default	single-host	
Configuration Mode	config interface ethernet	
History	3.4.2008	
	3.4.2300	Added “single-host” option
Example	`switch (config interface ethernet 1/2)# dot1x host-mode single-host`	
Related Commands		
Notes		

10.1.9.3.5 dot1x port-control

	dot1x port-control [auto | force-authorized | force-unauthorized] no dot1x port-control Configures 802.1x port access entity (PAE) port-control. The no form of the command resets the parameter to its default.	
Syntax Description	auto	The authenticator uses PAE authentication services to allow or block the port traffic
	force-authorized	Allows traffic on this port regardless of supplicant authorization
	force-unauthorized	Blocks traffic on this port regardless of supplicant authorization
Default	Force-authorized	
Configuration Mode	config interface ethernet	
History	3.4.2008	
Example	`switch (config interface ethernet 1/2)# dot1x port-control auto`	
Related Commands		

Notes	

10.1.9.3.6 dot1x radius-server host

	dot1x radius-server host <IP address> [enable | auth-port <port> | key <password> | prompt-key | retransmit <retries> | timeout <seconds>] no dot1x radius-server host <IP address> enable Configure 802.1x RADIUS server IP address. The no form of the command disables 802.1x RADIUS server.	
Syntax Description	auth-port	Sets 802.1x RADIUS port to use with this server Range: 1-65535
	enable	Sets 802.1x RADIUS as administratively enabled
	key	Configures 802.1x global RADIUS shared secret for servers
	prompt-key	Prompts for key, rather than entering on command line
	retransmit	Configure 802.1x global RADIUS retransmit count for servers Range: 0-5 seconds
	timeout	Configures 802.1x global RADIUS timeout value for servers Range: 1-60 seconds
Default	auth-port: 1812 key: empty string retransmit: 1 timeout: 3	
Configuration Mode	config	
History	3.4.2008	
Example	`switch (config)# dot1x radius-server host 10.10.10.10 auth-port 65535 prompt-key enable`	
Related Commands		
Notes	• The no form of the various parameters resets them to their default values as indicated in the Default section above • It is possible to configure up to 5 RADIUS servers • It is possible to configure only 1 authentication port per RADIUS server IP	

10.1.9.3.7 dot1x reauthenticate

	dot1x reauthenticate no dot1x reauthenticate Enables supplicant re-authentication according to the configuration of command "dot1x timeout reauthentication". The no form of the command disables supplicant re-authentication.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface ethernet
History	3.4.2008
Example	`switch (config interface ethernet 1/2)# dot1x reauthenticate`

Related Commands	
Notes	

10.1.9.3.8 dot1x system-auth-control

	dot1x system-auth-control no dot1x system-auth-control Enables the system as authenticator. The no form of the command disables the system as authenticator.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.4.2008
Example	`switch (config)# dot1x system-auth-control`
Related Commands	
Notes	

10.1.9.3.9 dot1x timeout reauthentication

	dot1x timeout reauthentication <period> no dot1x timeout reauthentication Configures the number of seconds between re-authentication attempts. The no form of the command resets the parameter to its default.	
Syntax Description	period	Time in second Range: 1-65535
Default	3600 seconds	
Configuration Mode	config interface ethernet	
History	3.4.2008	
Example	`switch (config interface ethernet 1/2)# dot1x timeout reauthentication 3600`	
Related Commands		
Notes		

10.1.9.3.10 dot1x timeout quiet-period

	dot1x timeout quiet-period <period> no dot1x timeout quiet-period Configures the number of seconds that the authenticator remains quiet following a failed authentication exchange with the supplicant. The no form of the command resets the parameter to its default.	
Syntax Description	period	Time in second Range: 1-65535
Default	60 seconds	
Configuration Mode	config interface ethernet	

History	3.4.2008
Example	`switch (config interface ethernet 1/2)# dot1x timeout quiet-period 60`
Related Commands	
Notes	

10.1.9.3.11 dot1x timeout tx-period

	dot1x timeout tx-period <period> no dot1x timeout tx-period Configures the maximum number of seconds that the authenticator waits for supplicant response of EAP-request/identify frame before retransmitting the request. The no form of the command resets the parameter to its default.	
Syntax Description	period	Time in second Range: 1-65535
Default	30 seconds	
Configuration Mode	config interface ethernet	
History	3.4.2008	
Example	`switch (config interface ethernet 1/2)# dot1x timeout quiet-period 30`	
Related Commands		
Notes		

10.1.9.3.12 dot1x max-req

	dot1x max-req <retries> no dot1x max-req Configures the maximum amount of retries for the authenticator to communicate with the supplicant over EAP. The no form of the command resets the parameter to its default.	
Syntax Description	retries	The number of request retries Range: 1-10
Default	2	
Configuration Mode	config interface ethernet	
History	3.4.2008	
Example	`switch (config interface ethernet 1/2)# dot1x max-req 2`	
Related Commands		
Notes		

10.1.9.3.13 show dot1x

	show dot1x Displays 802.1x information on all interfaces.
Syntax Description	N/A

Default	N/A
Configuration Mode	Any command mode
History	3.4.2008
Example	

```
switch (config)# show dot1x

System authentication is enabled
------------------------------------------------------------------------
Port        Pae         Host-mode   Port-control        Status
------------------------------------------------------------------------
Eth1/1      Enabled     multi-host  auto                unauthorized
Eth1/2      Disabled    multi-host  force-authorized    down
Eth1/3      Disabled    multi-host  force-authorized    down
Eth1/4      Disabled    multi-host  force-authorized    down
Eth1/5      Disabled    multi-host  force-authorized    down
Eth1/6      Disabled    multi-host  force-authorized    down
Eth1/7      Disabled    multi-host  force-authorized    down
Eth1/8      Disabled    multi-host  force-authorized    down
Eth1/9      Disabled    multi-host  force-authorized    down
...
```

Related Commands	
Notes	

10.1.9.3.14 show dot1x interfaces ethernet

	show dot1x interfaces ethernet <slot>/<port> Displays 802.1x interface information.	
Syntax Description	<slot>/<port>	Ethernet interface
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
Example		

```
switch (config)# show dot1x interfaces ethernet 1/2

Eth1/2
  PAE Status:                       Enabled
  Configured host mode:             Multi-host
  Configured port-control:          Auto
  Authentication status:            Unauthorized
  Re-Authentication:                Enabled
  Re-Authentication period (sec):   3600
  Tx wait period (sec):             30
  Quiet period (sec):               60
  Max request retry:                2
  Last EAPOL RX source MAC:         00:00:00:00:00:00
```

Related Commands	
Notes	

10.1.9.3.15 show dot1x interfaces ethernet statistics

	show dot1x interfaces ethernet <slot>/<port> statistics Displays 802.1x interface information.	
Syntax Description	<slot>/<port>	Ethernet interface
Default	N/A	

Configuration Mode	Any command mode
History	3.4.2008
Example	

```
switch (config)# show dot1x interfaces ethernet 1/2 statistics

Eth1/2
  EAPOL frames received:                    3
  EAPOL frames transmitted:                 2
  EAPOL Start frames received:              1
  EAPOL Logoff frames received:             0
  EAP Response-ID frames received:          2
  EAP Response frames received:             0
  EAP Request-ID frames transmitted:        2
  EAP Request frames transmitted:           0
  Invalid EAPOL frames received:            0
  EAP length error frames received:         0
  Last EAPOL frame version:                 1
  Last EAPOL frame source:                  00:1a:a0:02:e9:8e
```

Related Commands	
Notes	

10.1.9.3.16 show dot1x radius

	show dot1x radius Displays 802.1x RADIUS settings.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.4.2008

Example	`switch (config)# show dot1x radius`
	`802.1x RADIUS defaults:`
	`  Key:  ********`
	`  Timeout:  3`
	`  Retransmit:  1`
	`No 802.1x RADIUS servers configured.`

Related Commands	
Notes	

10.2 Cryptographic (X.509, IPSec) and Encryption

This page contains commands for configuring, generating and modifying x.509 certificates used in the system. Certificates are used for creating a trusted SSL connection to the system.

Crypto commands also cover IPSec configuration commands used for establishing a secure connection between hosts over IP layer which is useful for transferring sensitive information.

10.2.1 System File Encryption

This feature encrypts all sensitive data on NVIDIA systems including logs certificates, keys, etc.

To activate encryption on the switch:

1. Enable encryption and configure key location as USB (if you are using a USB device). Run:

```
switch (config)# crypto encrypt-data key-location usb key mypassword

Warning! All sensitive files are about to be encrypted
- System will perform reset factory, configuration files will be preserved
- System will be rebooted
- Active configuration will be preserved
- Do not power-off, wait for the system to boot

Type 'YES' to confirm this action: YES
```

> ***IMPORTANT***
> Encryption and decryption perform “reset factory keep-config” on the switch system once configured. This means that sysdumps, logs, and images are deleted.

> The key may be saved locally as well by using the parameter “local” instead of “usb” but that configuration is less secure.

2. After the system reboots, verify configuration. Run:

```
switch (config)# show crypto encrypt-data
Sensitive files encryption:
   Status:          enabled
   Key location:    usb
   Cipher:          aes256
```

> Once encryption is enabled, reverting back to an older version while encrypted is not possible. The command “no crypto encrypt-data” must be run before attempting to downgrade to an older OS version.

> If encryption is enabled, upgrading to a new OS version maintains the encryption configuration.

10.2.2 Cryptographic and Encryption Commands

- 10.2.1 System File Encryption
- 10.2.2 Cryptographic and Encryption Commands
  - 10.2.2.1 crypto encrypt-data
  - 10.2.2.2 crypto ipsec ike
  - 10.2.2.3 crypto ipsec peer local
  - 10.2.2.4 crypto certificate ca-list
  - 10.2.2.5 crypto certificate default-cert
  - 10.2.2.6 crypto certificate generation
  - 10.2.2.7 crypto certificate name
  - 10.2.2.8 crypto certificate system-self-signed
  - 10.2.2.9 show crypto certificate
  - 10.2.2.10 show crypto encrypt-data
  - 10.2.2.11 show crypto ipsec

10.2.2.1 crypto encrypt-data

	crypto encrypt-data key-location <local | usb> key <password> no crypto encrypt-data Enables and configures system file encryption. The no form of the command decrypts sensitive information on the system.	
Syntax Description	key-location	Configures where to store the encryption key: • local—stores the key locally • usb—stores the key on a USB device
	key	Configures a key
Default	N/A	
Configuration Mode	config	
History	3.6.1002	
Example		
Related Commands	show crypto certificate	
Notes	• It is recommended to store the encryption password on a USB device rather than locally • Enabling encryption may slightly slow system performance • If the key is stored on the USB, it must be plugged into the switch in order for the switch to boot. After the switch has booted, the USB key is no longer required and, for security purposes, it is recommended to remove it after running “usb eject”. The USB key may be needed again if the switch is rebooted or if the switch needs to be decrypted.	

10.2.2.2 crypto ipsec ike

	crypto ipsec ike {clear sa [peer {any | <IPv4 or IPv6 address>} local <IPv4 or IPv6 address>] | restart} Manages the IKE (ISAKMP) process or database state.	
Syntax Description	clear	Clears IKE (ISAKMP) peering state
	sa	Clears IKE generated ISAKMP and IPSec security associations (remote peers are affected)
	peer	Clears security associations for the specified IKE peer (remote peers are affected) • all—clears security associations for all IKE peerings with a specific local address (remote peers are affected) • IPv4 or IPv6 address—clears security associations for specific IKE peering with a specific local address (remote peers are affected)
	IPv4 or IPv6 address	Clears security associations for the specified IKE peering (remote peer is affected)
	local	Clear security associations for the specified/all IKE peering (remote peer is affected)
	restart	Restarts the IKE (ISAKMP) daemon (clears all IKE state, peers may be affected)
Default	N/A	
Configuration Mode	config	

History	3.2.3000
Example	`switch (config)# crypto ipsec ike restart`
Related Commands	show crypto certificate
Notes	

10.2.2.3 crypto ipsec peer local

	crypto ipsec peer local {enable | keying {ike negotiation {ikev1 | ikev2} | [auth { hmac-sha1 | hmac-sha256 | hmac-sha512 | aes-xcbc} | dh-group | disable | encrypt { 3des-cbc| aes-cbc | aes-gcm} | exchange-mode | lifetime | local | mode | peer-identity | pfs-group | preshared-key | prompt-preshared-key | transform-set] | manual [auth | disable | encrypt | local-spi | mode | remote-spi]}} Configures IPSec in the system.	
Syntax Description	enable	Enables IPSec peering.
	ike	Configures IPSec peering using IKE ISAKMP to manage SA keys. The following optional parameters are available: • auth—configures the authentication algorithm for IPSec peering • dh-group—configures the phase1 Diffie-Hellman group proposed for secure IKE key exchange • disable—configures this IPSec peering administratively disabled • encrypt—configures the encryption algorithm for IPSec peering • exchange-mode—configures the IKE key exchange mode to propose for peering • lifetime—configures the SA lifetime to propose for this IPSec peering • local-identity—configures the ISAKMP payload identification value to send as local endpoint's identity • mode—configures the peering mode for this IPSec peering • peer-identity—configures the identification value to match against the peer's ISAKMP payload identification • pfs-group—configures the phase2 PFS (Perfect Forwarding Secrecy) group to propose for Diffie-Hellman exchange for this IPSec peering • preshared-key—configures the IKE pre-shared key for the IPSec peering • prompt-preshared-key—prompts for the pre-shared key, rather than entering it on the command line • transform-set—configures transform proposal parameters
	keying	Configures key management for this IPSec peering. • auth—configures the authentication algorithm for this IPSec peering • disable—configures this IPSec peering administratively disabled • encrypt—configures the encryption algorithm for this IPSec peering • local-spi—configures the local SPI for this manual IPSec peering • mode—configures the peering mode for this IPSec peering • remote-spi—configures the remote SPI for this manual IPSec peering
	manual	Configures IPSec peering using manual keys.

Default	N/A
Configuration Mode	config
History	3.2.3000 3.9.3100: Added support for IKEv2 and new ciphers
Example	`switch (config)# crypto ipsec peer 10.10.10.10 local 10.7.34.139 enable`
Related Commands	show crypto certificate
Notes	As of version 3.9.3100, NULL will not be supported as an authentication or encryption algorithm for IPsec peering. New ciphers are supported (hmac-sha512 and aes-xcbc for authentication and aes-gcm for encryption. 1, 2, 5, 22, 23, 24 pfs/dh-groups will not be supported, while 19, 20, 21 will be supported only with IKEv2. The transform-set options ah-and-esp-ah are no longer supported. Libreswan is used instead of openswan.

10.2.2.4 crypto certificate ca-list

	crypto certificate ca-list [default-ca-list name {<cert-name> | system-self-signed}] no crypto certificate ca-list [default-ca-list name {<cert-name> | system-self-signed}] Adds the specified CA certificate to the default CA certificate list. The no form of the command removes the certificate from the default CA certificate list.	
Syntax Description	cert-name	The name of the certificate
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Example	`switch (config) # crypto certificate default-cert name test`	
Related Commands	show crypto certificate	
Notes	• Two certificates with the same subject and issuer fields cannot both be placed onto the CA list • The no form of the command does not delete the certificate from the certificate database • Unless specified otherwise, applications that use CA certificates will still consult the well-known certificate bundle before looking at the default-ca-list	

10.2.2.5 crypto certificate default-cert

	crypto certificate default-cert name {<cert-name> | system-self-signed} no crypto certificate default-cert name {<cert-name> | system-self-signed} Designates the named certificate as the global default certificate role for authentication of this system to clients. The no form of the command reverts the default-cert name to “system-self-signed” (the “cert-name” value is optional and ignored).	
Syntax Description	cert-name	The name of the certificate
Default	N/A	
Configuration Mode	config	

History	3.2.3000
Example	`switch (config) # crypto certificate default-cert name test`
Related Commands	show crypto certificate
Notes	• A certificate must already be defined before it can be configured in the default-cert role • If the named default-cert is deleted from the database, the default-cert automatically becomes reconfigured to the factory default, the “system-self-signed” certificate

10.2.2.6 crypto certificate generation

	crypto certificate generation default {country-code | days-valid > | ca-valid <true/false> | email-addr | hash-algorithm {sha1 | sha256} | key-size-bits | locality | org-unit | organization | state-or-prov} Configures default values for certificate generation.	
Syntax Description	country-code	Configures the default certificate value for country code with a two-alphanumeric-character code or -- for none.
	days-valid	Configures the default certificate valid days Default value: 365 days
	email-addr	Configures the default certificate value for email address
	hash-algorithm {sha1 | sha256}	Configures the default certificate hashing algorithm
	key-size-bits	Configures the default certificate value for private key size (private key length in bits—at least 1024, but 2048 is strongly recommended)
	locality	Configures the default certificate value for locality
	org-unit	Configures the default certificate value for organizational unit
	organization	Configures the default certificate value for the organization name
	state-or-prov	Configures the default certificate value for state or province
	ca-valid {true | false}	Configures the default certificate CA Basic Constraints flag set to TRUE/FALSE
Default	hash-algorithm - sha1	
Configuration Mode	config	
History	3.2.1000 3.3.4350: Added “hash-algorithm” parameter 3.6.4000: Added “days-valid” parameter 3.8.2100: Added "ca-valid" parameter	
Example	`switch (config) # crypto certificate generation default hash-algorithm sha256`	
Related Commands	show crypto certificate	
Notes		

10.2.2.7 crypto certificate name

	crypto certificate name {<cert-name> | system-self-signed} {comment <new comment> | generate selfsigned [comment <cert-comment> | common-name <domain> | country-code <code> | days-valid <days> | ca-valid <true/false> | email-addr <address> | hash-algorithm {sha1 | sha256} | key-size-bits <bits> | locality <name> | org-unit <name> | organization <name> | serial-num <number> | state-or-prov <name>]} | private-key pem <PEM string> | prompt-private-key | public-cert [comment <comment string> | pem <PEM string>] | regenerate days-valid <days> | ca-valid <true/false> | rename <new name>} no crypto certificate name <cert-name> Configures default values for certificate generation. The no form of the command clears/deletes certain certificate settings.	
Syntax Description	cert-name	Unique name by which the certificate is identified.
	comment	Specifies a certificate comment.
	generate self-signed	Generates certificates. This option has the following parameters which may be entered sequentially in any order: • comment—specifies a certificate comment (free string) • common-name—specifies the common name of the issuer and subject (e.g. a domain name) • country-code—specifies the country codwo-alphanumeric-character country code, or “--” for none) • days-valid—specifies the number of days the certificate is valid • email-addr—specifies the email address • hash-algorithm—specifies the hashing function used for signature algorithm. Default value is SHA256. • key-size-bits—specifies the size of the private key in bits (private key length in bits - at least 1024 but 2048 is strongly recommended) • locality—specifies the locality name • org-unit—specifies the organizational unit name • organization—specifies the organization name • serial-num—specifies the serial number for the certificate (a lower-case hexadecimal serial number prefixed with “0x”) • state-or-prov—specifies the state or province name • ca-valid—Specifies certificate CA Basic Constraints flag set to TRUE/FALSE
	private-key pem	Specifies certificate contents in PEM format
	prompt-private-key	Prompts for certificate private key with secure echo
	public-cert	Installs a certificate
	regenerate	Regenerates the named certificate using configured certificate generation default values for the specified validity period
	rename	Renames the certificate
Default	N/A	
Configuration Mode	config	
History	3.2.3000 3.3.4402: Added “hash-algorithm” parameter 3.6.4000: Added “days-valid” parameter 3.8.2100: Added "ca-valid" parameter	

Example	`switch (config) # crypto certificate name system-self-signed generate self-signed hash-algorithm sha256`
Related Commands	show crypto certificate
Notes	

10.2.2.8 crypto certificate system-self-signed

	crypto certificate system-self-signed regenerate {[days-valid <days>] | ca-valid <true/false>} Configures default values for certificate generation.	
Syntax Description	days-valid	Specifies the number of days the certificate is valid
	ca-valid	Specifies certificate CA Basic Constraints flag set to TRUE/FALSE
Default	N/A	
Configuration Mode	config	
History	3.2.1000 3.8.2100: Added the ca-valid option	
Example	`switch (config) # crypto certificate system-self-signed regenerate days-valid 3` `switch (config) # crypto certificate system-self-signed regenerate ca-valid false`	
Related Commands	show crypto certificate	
Notes		

10.2.2.9 show crypto certificate

	show crypto certificate [detail | public-pem | default-cert [detail | public-pem] | [name <cert-name> [detail | public-pem] | ca-list [default-ca-list]] Displays information about all certificates in the certificate database.	
Syntax Description	ca-list	Displays the list of supplemental certificates configured for the global default system CA certificate role
	default-ca-list	Displays information about the currently configured default certificates of the CA list
	default-cert	Displays information about the currently configured default certificate
	detail	Displays all attributes related to the certificate
	name	Displays information about the certificate specified
	public-pem	Displays the uninterpreted public certificate as a PEM formatted data string
Default	N/A	
Configuration Mode	config	
History	3.2.1000 3.8.2100: Updated output	
Example		

```
switch (config) # show crypto certificate

Certificate with name 'system-self-signed' (default-cert)
    Comment:                    system-generated self-signed certificate
    Private Key:                present
    Serial Number:              0x546c935511bcafc21ac0e8249fbe0844
    SHA-1 Fingerprint:          fe6df38dd26801971cb2d44f62dbe492b6063c5f

    Validity:
        Starts:                 2012/12/02 13:45:05
        Expires:                2013/12/02 13:45:05

    Subject:
        Common Name:            IBM-DEV-Bay4
        Country:                IS
        State or Province:
        Locality:
        Organization:
        Organizational Unit:
        E-mail Address:

    Issuer:
        Common Name:            IBM-DEV-Bay4
        Country:                IS
        State or Province:
        Locality:
        Organization:
        Organizational Unit:
        E-mail Address:

    X509 Extensions:
        Basic Constraints:
            CA: TRUE
```

Related Commands	
Notes	

10.2.2.10 show crypto encrypt-data

	show encrypt-data Displays sensitive data encryption information.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.1002

Example	`switch (config)# show crypto encrypt-data`
	`Sensitive files encryption:`
	`  Status:  enabled`
	`  Key location:  usb`
	`  Cipher:  aes256`

Related Commands	
Notes	

10.2.2.11 show crypto ipsec

	show crypto ipsec [brief | configured | ike | policy | sa] Displays information ipsec configuration.
Syntax Description	N/A
Default	N/A

Configuration Mode	config
History	3.2.1000
Example	switch (config)# show crypto ipsec IPSec Summary ------------- Crypto IKE is using pluto (Openswan) daemon. Daemon process state is stopped. No IPSec peers configured. IPSec IKE Peering State ---------------------- Crypto IKE is using pluto (Openswan) daemon. Daemon process state is stopped. No active IPSec IKE peers. IPSec Policy State ------------------ No active IPSec policies. IPSec Security Association State -------------------------------- No active IPSec security associations.
Related Commands	
Notes	

11 Quality of Service (QoS)

11.1 QoS Classification

QoS classification assigns a QoS class to the packet. The QoS class of the packet is indicated internally in the switch using the switch-priority parameter (8 possible values).

Switch-priority affects the packet buffering and transmission scheduling. There are 8 possible values for switch-priority. The classification is based on the PCP and DEI fields in the VLAN tag, the DSCP field in the IP header. In addition, the default value can be configured for the incoming port. And the switch-priority of the packet also can be reconfigured by the ACL.

The switch-priority of the packet is used for priority fields re-marking at the egress.

11.1.1 Trust Levels

QoS classification depends on the port configuration for QoS trust level which determines which packet header fields derive the switch-priority. The following trust states are supported:

- Trust port
  - Based on port default settings
- Trust L2 (PCP,DEI)
  - Based on packet PCP,DEI fields for VLAN tagged packets
  - Else, based on the port default setting for VLAN un-tagged packets
- Trust L3 (DSCP)
  - Based on packet DSCP field for IP packets
  - Else, based on port default setting for non-IP
- Trust both
  - Based on packet DSCP for IP packets
  - Else, based on packet PCP,DEI for VLAN tagged packets
  - Else, based on the port default setting

The following table and figure summarize the packet classification rules.

Packet Type		QoS Classification Config (per Interface)			
IP/MPLS	VLAN	Trust Both	Trust L3	Trust L2	Trust Port
IP/MPLS	Tagged	DSCP	DSCP	PCP,DEI	Port Default
IP/MPLS	Untagged	DSCP	DSCP	Port Default	Port Default
non-IP/MPLS	Tagged	PCP,DEI	Port Default	PCP,DEI	Port Default
non-IP/MPLS	Untagged	Port Default	Port Default	Port Default	Port Default

Default switch-priority is configured as trust L2.

11.1.2 Switch Priority to IEEE Priority Mapping

IEEE defines priority value for a packet which is used in the switch for the pause flow control.

The device maps the switch-priority into IEEE priority value using device global switch priority to IEEE priority table.

11.1.3 Default QoS Configuration

Parameter	Range	Configuration
Trust level	All ports	Trust L2
DSCP to switch-priority	0-7	0
DSCP to switch-priority	8-15	1
DSCP to switch-priority	16-23	2
DSCP to switch-priority	24-31	3
DSCP to switch-priority	32-39	4
DSCP to switch-priority	40-47	5
DSCP to switch-priority	48-55	6
DSCP to switch-priority	56-63	7
PCP to switch-priority	0	0
PCP to switch-priority	1	1
PCP to switch-priority	2	2
PCP to switch-priority	3	3
PCP to switch-priority	4	4
PCP to switch-priority	5	5
PCP to switch-priority	6	6
PCP to switch-priority	7	7
Port PCP,DEI default	All ports	0
Port switch-priority when “trust port” is enabled	All ports	0
Switch-priority to IEEE priority	0	0
Switch-priority to IEEE priority	1	1
Switch-priority to IEEE priority	2	2
Switch-priority to IEEE priority	3	3
Switch-priority to IEEE priority	4	4
Switch-priority to IEEE priority	5	5
Switch-priority to IEEE priority	6	6
Switch-priority to IEEE priority	7	7

11.1.4 Control Protocols

Protocol	Switch Priority
xSTP	Switch Priority 7

Protocol	Switch Priority
LACP	Switch Priority 7
LLDP	Switch Priority 7
PTP	Interface VLAN: Switch Priority 7
	Router Port: Switch Priority 6
BGP	Switch Priority 6
OSPF	Switch Priority 6
PIM	Switch Priority 6
IGMP	Switch Priority 6
MLAG	Switch Priority 6
SFLOW	Switch Priority 6
VRRP	Switch Priority 6

11.2 QoS Rewrite

NVIDIA Spectrum enables rewriting QoS identifier values (DSCP, PCP, DEI) of incoming packets.

The configuration for preserving the values or rewriting them is set per ingress port. The configuration of the new values is set per egress port and is based on the mapping from the switch-priority.

In addition, the packets that pass the router module in the switch can be configured to change the “rewrite enable” configuration as well as the switch-priority.

11.2.1 Switch-priority to PCP,DEI Re-marking Mapping

Packet PCP and DEI fields can be updated by the switch based on switch-priority to PCP,DEI mapping tables. The mapping can be configured per egress port.
The reason for the mapping is to enable changing interpretation between two administrative domains in the network, or when a source of data is not fully trusted, and the default values are not desired. This mapping takes effect after deriving switch-priority from the PCP,DEI fields.

11.2.2 Switch-priority to DSCP Re-marking Mapping

Packet DSCP field can be updated based on switch-priority to DSCP mapping tables. The mapping can be configured per egress port. MPLS packets are untouched regardless this setting.

The reason for the mapping is to enable changing interpretation between two administrative domains in the network, or when a source of data is not fully trusted. This mapping will take effect after deriving switch-priority from the DSCP field.

11.2.3 DSCP to Switch-priority in Router

Spectrum enables mapping of DSCP to switch-priority in the router using a global mapping table. This mapping has global configuration for whether to change the “Rewrite/Preserve PCP,DEI” bit. This configuration sets how the DSCP to switch-priority would affect the packet.

11.2.4 Default Configuration

- By default no ingress rewrite configuration is set
- By default PCP rewrite configuration in router is set
- The default mapping is as following:
    - Switch-priority=i to PCP,DEI=i,0, i=0-7
    - Switch-priority=i to DSCP=8i, i=0-7

11.3 Queuing and Scheduling (ETS)

Enhanced Transmission Selection (ETC) provides a common management framework for assignment of bandwidth to traffic classes, for weighted round robin (WRR) scheduling. If a traffic class does not use all the bandwidth allocated to it, other traffic classes can use the available bandwidth. This allows optimal utilization of the network capacity while prioritizing and providing the necessary resources.

The ETS feature has the following attributes:

- ETS global admin
    - Enable (default)—scheduling mode is WRR according to the configured bandwidth-per-traffic class
    - Disable-scheduling mode is Strict Priority (SP)
- Bandwidth percentage for each traffic class: by default each traffic class gets an equal share

After the output port of the packet is determined and the packet is buffered, it is queued for transmission. Each egress port is combined from the multi-level queuing structure. The scheduling of transmission from the queues relies on various configurations such as ETS weight, flow control, rate shaping etc.

11.3.1 Traffic Class

The switch-priority of the packet assigns it to a specific traffic class (TClass). The TClass of the packet determines the packet path in the queuing structure. There are 8 TCs supported by the system.

11.3.2 Traffic Shapers

11.3.2.1 Maximum Shapers

TCs can be configured for rate shaping as described in the following:

- TClass queues: shaper per TClass queue

- Port: shaper per port (bytes only)

Shapers support the following configurations:

- Committed Incoming Rate (CIR) [bits/packets per second]
- Committed Burst Size (CBS) [bits/packets]

Each shaper has granularity rate of 1Mb/s, 10Mb/s, 100Mb/s and 1Gb/s (or 128K, 1280K, 12M, 128M pps). The maximum CBS is 3GB or 384M packets.

11.3.2.2 Minimum Shapers

TC queues can be configured for minimal rate shaping. The minimum shaper configuration overrides all other scheduling configurations. So that if ETS or WRR scheduling allocates to a TC queue lower rate than the configured minimum, that queue receives strictly higher priority over the others. If several queues receive a rate below the configured minimum, the arbitration between them can be configured as a WRR, or as strict according to the queue index.

The configuration of min shaper is identical to the configuration of max shaper.

11.3.3 Default Shaper Configuration

Parameter	Range	Configuration
Switch-priority to TC	0	0
Switch-priority to TC	1	1
Switch-priority to TC	2	2
Switch-priority to TC	3	3
Switch-priority to TC	4	4
Switch-priority to TC	5	5
Switch-priority to TC	6	6
Switch-priority to TC	7	7
Shaping	All ports	No max/min shaping configured

11.4 RED and ECN

Random early detection (RED) is a mechanism that randomly drops packets before the switch buffer fills up in case of congestion. Explicit congestion notification (ECN) is used for congestion control protocols (TCP and RoCE CC – DCQCN) to handle congestion before packets are dropped. RED and ECN can be configured separately or concurrently per traffic class.

Relative RED/ECN is supported on TC queues. This allows the thresholds of the drop/mark actions to behave relatively to the dynamic thresholds configured for the shared buffer.

RED/ECN drop profiles are defined according to 2 parameters as shown in the following figure:

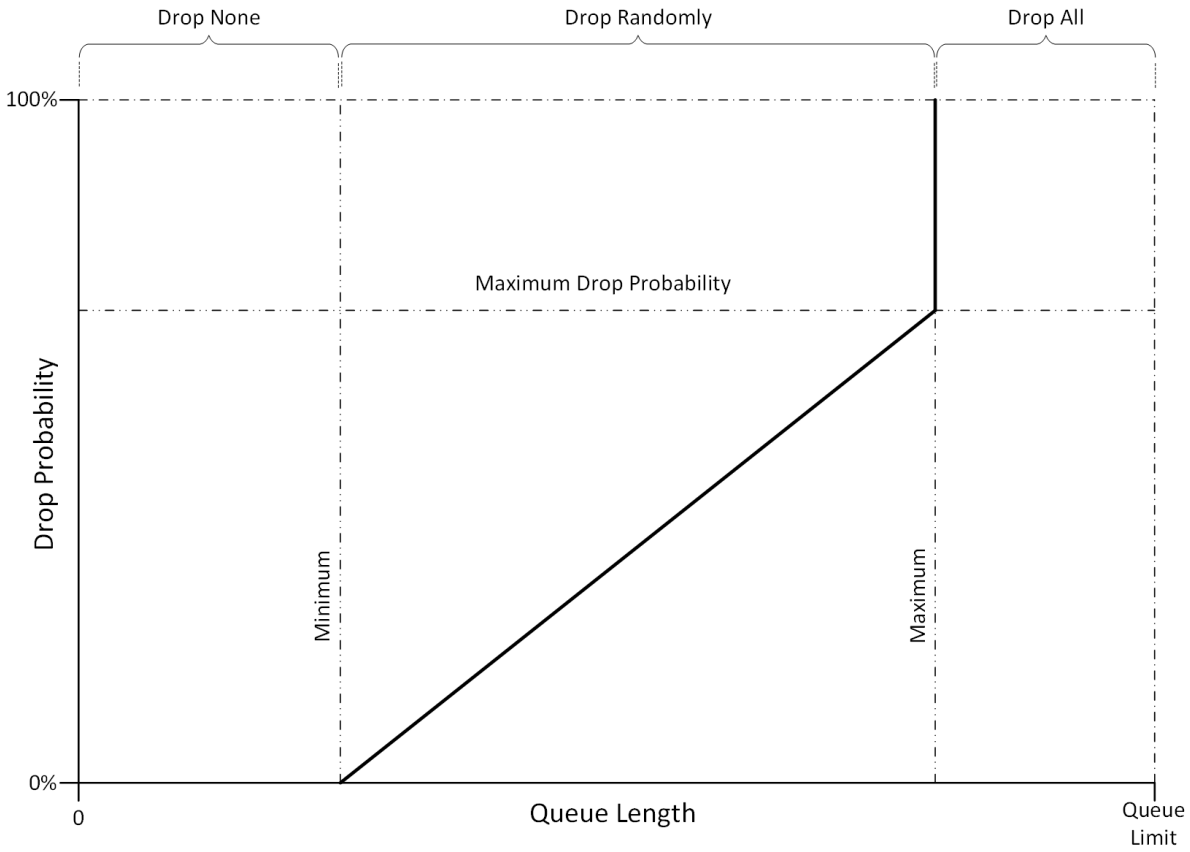

- Minimum – a threshold that defines the average queue length below which the packets are not dropped/marked
- Maximum – a threshold that defines the average queue length above which the packets are always dropped/marked

It is possible to configure the minimum and maximum thresholds to have the same value which would represent a step function from “drop none” to “drop all”.

> RED/ECN is only supported for unicast traffic classes.

11.5 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community posts:

- End-to-End QoS Configuration for Switches (SwitchX) and Adapters
- How To Configure DSCP-Based PFC on Spectrum Switches
- HowTo Enable PFC on Switches (SwitchX)
- HowTo Configure QoS on Switches (SwitchX)
- Understanding TC Scheduling on Spectrum Switches (WRR, SP)
- HowTo Configure ECN on Ethernet Switches (Spectrum)
- Understanding QoS Classification (Trust) on Spectrum Switches
- QoS Tuning on Spectrum Switches – FAQ

11.6 QoS Commands

- QoS Commands
- Priority Flow Control (PFC)
- Shared Buffers
- Storm Control
- Head-of-Queue Lifetime Limit
- Store-and-Forward

11.7 QoS Commands

- 11.7.1 QoS Classification
  - 11.7.1.1 vlan default priority
  - 11.7.1.2 vlan default dei
  - 11.7.1.3 qos trust
  - 11.7.1.4 qos default switch-priority
  - 11.7.1.5 qos map pcp dei
  - 11.7.1.6 qos map dscp
  - 11.7.1.7 show interfaces ethernet counters pfc prio
  - 11.7.1.8 show qos
  - 11.7.1.9 show qos interface ethernet
  - 11.7.1.10 show qos interface mlag-port-channel
  - 11.7.1.11 show qos interface port-channel
  - 11.7.1.12 show qos interface l2-mapping
  - 11.7.1.13 show qos interface l3-mapping
  - 11.7.1.14 show qos interface rewrite-mapping
  - 11.7.1.15 show qos interface tc-mapping
  - 11.7.1.16 show qos mapping ingress interface egress interface
- 11.7.2 QoS Rewrite
  - 11.7.2.1 qos rewrite pcp
  - 11.7.2.2 qos rewrite dscp
  - 11.7.2.3 qos rewrite map switch-priority pcp dei
  - 11.7.2.4 qos rewrite map switch-priority dscp
  - 11.7.2.5 qos ip rewrite pcp
  - 11.7.2.6 show qos ip rewrite
- 11.7.3 Queuing and Scheduling (ETS)
  - 11.7.3.1 bind switch-priority
  - 11.7.3.2 bandwidth guaranteed
  - 11.7.3.3 bandwidth shape
  - 11.7.3.4 show dcb ets
- 11.7.4 RED & ECN
  - 11.7.4.1 traffic-class congestion-control
  - 11.7.4.2 show interfaces ethernet congestion-control

11.7.1 QoS Classification

11.7.1.1 vlan default priority

	vlan default priority [<priority>] no vlan default priority [<priority>] Configures default PCP for packets arrived without VLAN tag. The no form of the command resets the value to its default.	
Syntax Description	priority	Range: 0-7
Default	0	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.1002	
Example	`switch (config interface ethernet 1/1) # vlan default priority 0`	
Related Commands		
Notes		

11.7.1.2 vlan default dei

	vlan default dei [<dei>] no vlan default dei [<dei>] Configures default DEI for packets arrived without VLAN tag. The no form of the command resets the value to its default.
Syntax Description	N/A
Default	0
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.6.1002
Example	`switch (config interface ethernet 1/1) # vlan default dei 0`
Related Commands	
Notes	

11.7.1.3 qos trust

	qos trust [port | L2 | L3 | both] no qos trust Configures QoS trust mode for the interface. The no form of the command resets the value to its default.
Syntax Description	N/A
Default	L2

Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.1002	
	3.8.1000	Updated notes
Example	`switch (config interface ethernet 1/1) # qos trust L3`	
Related Commands		
Notes	Please see the table presenting packet classification rules for more information	

11.7.1.4 qos default switch-priority

	qos default switch-priority [<switch-priority>] no qos default switch-priority [<switch-priority>] Configures default switch-priority for the interface when “port” trust mode is active, or for non-IP and untagged packets in other trust modes. The no form of the command resets the value to its default.	
Syntax Description	switch-priority	Range: 0-7
Default	0	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.1002	
	3.7.0000	Edited command definition
Example	`switch (config interface ethernet 1/1) # qos default switch-priority 0`	
Related Commands	qos trust	
Notes		

11.7.1.5 qos map pcp dei

	qos map pcp <0-7> dei <0-1> to switch-priority <0-7> Configures interface PCP, DEI to switch-priority mapping for IP/MPLS and non-IP/MPLS tagged packets in “L2” trust mode and for non-IP/MPLS tagged packets in “both” trust mode. The no form of the command resets the value to its default.
Syntax Description	N/A
Default	PCP to switch-priority mapping: 0 → 0 1 → 1 2 → 2 3 → 3 4 → 4 5 → 5 6 → 6 7 → 7

Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.1002	
	3.8.2100	Updated example
Example	`switch (config interface ethernet 1/1) # qos map pcp 5 dei 1 to switch-priority 7`	
Related Commands	qos trust	
Notes		

11.7.1.6 qos map dscp

	qos map dscp <dscp> [to switch-priority <switch-priority>] no qos map dscp <dscp> [to switch-priority <switch-priority>] Configures interface DSCP to switch-priority mapping in “L3” or “both” trust mode. The no form of the command resets the value to its default.	
Syntax Description	switch-priority	Range: 0-7
	dscp	Range: 0-63
Default	DSCP to switch-priority mapping:	0-7 → 0 8-15 → 1 16-23 → 2 24-31 → 3 32-39 → 4 40-47 → 5 48-55 → 6 56-63 → 7
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.1002	
Example	`switch (config interface ethernet 1/1) # qos map dscp 45`	
Related Commands	qos trust	
Notes		

11.7.1.7 show interfaces ethernet counters pfc prio

	show interfaces ethernet [<slot/port> | <slot/port>-<slot/port>] counters pfc prio <priority> Displays priority flow control counters for the specified interface and priority.	
Syntax Description	slot/port	Number of Ethernet interface in form of slot/port
	priority	Valid priority values: 0-7 or all
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
	3.9.1000	Added ability to use a range of ports

Example	switch (config) # show interfaces ethernet 1/1-1/2 counters pfc prio 1 Eth1/1: PFC 1: Rx: 0 pause packets 0 pause duration Tx: 0 pause packets 0 pause duration Eth1/2: PFC 1: Rx: 0 pause packets 0 pause duration Tx: 0 pause packets 0 pause duration
Related Commands	
Notes	From version 3.9.1000 and up, the "slot/port" attribute is optional. If nothing is selected, information for all ports will be displayed

11.7.1.8 show qos

	show qos Displays QoS information.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	
	3.6.8008	Updated example
Example		

```
switch (config) # show qos

Eth1/1:
  Trust mode             : L2
  Default switch-priority: 0
  Default PCP            : 0
  Default DEI            : 0
  PCP,DEI rewrite        : disabled
  IP PCP;DEI rewrite     : enable
  DSCP rewrite           : disabled

  PCP(DEI); DSCP to switch-priority mapping:
    ------------------------------------------------------------------
    PCP(DEI)              DSCP                      switch-priority
    ------------------------------------------------------------------
    0(0) 0(1)             0 1 2 3 4 5 6 7           0
    1(0) 1(1)             8 9 10 11 12 13 14 15     1
    2(0) 2(1)             16 17 18 19 20 21 22 23   2
    3(0) 3(1)             24 25 26 27 28 29 30 31   3
    4(0) 4(1)             32 33 34 35 36 37 38 39   4
    5(0) 5(1)             40 41 42 43 44 45 46 47   5
    6(0) 6(1)             48 49 50 51 52 53 54 55   6
    7(0) 7(1)             56 57 58 59 60 61 62 63   7

  PCP(DEI); DSCP rewrite mapping (switch-priority to PCP(DEI); DSCP; traffic-class):
    Egress Interface: Eth1/1

    ----------------------------------------
    switch-priority   PCP(DEI)   DSCP     TC
    ----------------------------------------
    0                 0(0)       0        0
    1                 1(0)       8        1
    2                 2(0)       16       2
    3                 3(0)       24       3
    4                 4(0)       32       4
    5                 5(0)       40       5
    6                 6(0)       48       6
    7                 7(0)       56       7
...
```

Related Commands	
Notes	

11.7.1.9 show qos interface ethernet

	show qos interface ethernet <port-id> Display QoS information for Ethernet interface.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
	3.6.8008	Updated example
Example		

```
switch (config)# show qos interface ethernet 1/1

Eth1/1:
  Trust mode             : L2
  Default switch-priority: 0
  Default PCP            : 0
  Default DEI            : 0
  PCP,DEI rewrite        : disabled
  IP PCP;DEI rewrite     : enable
  DSCP rewrite           : disabled

  PCP(DEI); DSCP to switch-priority mapping:
    ------------------------------------------------------------------
    PCP(DEI)              DSCP                      switch-priority
    ------------------------------------------------------------------
    0(0) 0(1)             0 1 2 3 4 5 6 7           0
    1(0) 1(1)             8 9 10 11 12 13 14 15     1
    2(0) 2(1)             16 17 18 19 20 21 22 23   2
    3(0) 3(1)             24 25 26 27 28 29 30 31   3
    4(0) 4(1)             32 33 34 35 36 37 38 39   4
    5(0) 5(1)             40 41 42 43 44 45 46 47   5
    6(0) 6(1)             48 49 50 51 52 53 54 55   6
    7(0) 7(1)             56 57 58 59 60 61 62 63   7

  PCP(DEI); DSCP rewrite mapping (switch-priority to PCP(DEI); DSCP; traffic-class):
    Egress Interface: Eth1/1

    ----------------------------------------
    switch-priority   PCP(DEI)   DSCP     TC
    ----------------------------------------
    0                 0(0)       0        0
    1                 1(0)       8        1
    2                 2(0)       16       2
    3                 3(0)       24       3
    4                 4(0)       32       4
    5                 5(0)       40       5
    6                 6(0)       48       6
    7                 7(0)       56       7
```

Related Commands	
Notes	

11.7.1.10 show qos interface mlag-port-channel

	show qos interface mlag-port-channel <port-id> Display QoS information for MPO.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
	3.6.6000	Updated example
Example		

```
switch (config)# show qos interface mlag-port-channel 1

Mpo1
 Trust mode: L2
 Default switch-priority: 0
 Default PCP: 0
 Default DEI: 0
 PCP,DEI rewrite: disabled
 IP PCP;DEI rewrite: enable
 DSCP rewrite: disabled

 PCP(DEI); DSCP to switch-priority mapping:
-------------------------------------------------------------------
 PCP(DEI)              DSCP                      switch-priority
-------------------------------------------------------------------
 0(0) 0(1)             0 1 2 3 4 5 6 7           0
 1(0) 1(1)             8 9 10 11 12 13 14 15     1
 2(0) 2(1)             16 17 18 19 20 21 22 23   2
 3(0) 3(1)             24 25 26 27 28 29 30 31   3
 4(0) 4(1)             32 33 34 35 36 37 38 39   4
 5(0) 5(1)             40 41 42 43 44 45 46 47   5
 6(0) 6(1)             48 49 50 51 52 53 54 55   6
 7(0) 7(1)             56 57 58 59 60 61 62 63   7

 PCP(DEI); DSCP rewrite mapping (switch-priority to PCP(DEI); DSCP; traffic-class):

 Egress Interface: Mpo1
-----------------------------------------
 switch-priority   PCP(DEI)   DSCP     TC
-----------------------------------------
 0                 0(0)       0        0
 1                 1(0)       8        1
 2                 2(0)       16       2
 3                 3(0)       24       3
 4                 4(0)       32       4
 5                 5(0)       40       5
 6                 6(0)       48       6
 7                 7(0)       56       7
```

Related Commands	
Notes	

11.7.1.11 show qos interface port-channel

	show qos interface port-channel <port-id> Display QoS information for port-channel interface.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
	3.6.8008	Updated example
Example		

```
switch (config)# show qos interface port-channel 1

Po1:
  Trust mode             : L2
  Default switch-priority: 0
  Default PCP            : 0
  Default DEI            : 0
  PCP,DEI rewrite        : disabled
  IP PCP;DEI rewrite     : enable
  DSCP rewrite           : disabled

  PCP(DEI); DSCP to switch-priority mapping:
    ----------------------------------------------------------------
    PCP(DEI)              DSCP                      switch-priority
    ----------------------------------------------------------------
    0(0) 0(1)             0 1 2 3 4 5 6 7           0
    1(0) 1(1)             8 9 10 11 12 13 14 15     1
    2(0) 2(1)             16 17 18 19 20 21 22 23   2
    3(0) 3(1)             24 25 26 27 28 29 30 31   3
    4(0) 4(1)             32 33 34 35 36 37 38 39   4
    5(0) 5(1)             40 41 42 43 44 45 46 47   5
    6(0) 6(1)             48 49 50 51 52 53 54 55   6
    7(0) 7(1)             56 57 58 59 60 61 62 63   7

  PCP(DEI); DSCP rewrite mapping (switch-priority to PCP(DEI); DSCP; traffic-class):
    Egress Interface: Po1

    ----------------------------------------
    switch-priority   PCP(DEI)   DSCP     TC
    ----------------------------------------
    0                 0(0)       0        0
    1                 1(0)       8        1
    2                 2(0)       16       2
    3                 3(0)       24       3
    4                 4(0)       32       4
    5                 5(0)       40       5
    6                 6(0)       48       6
    7                 7(0)       56       7
```

Related Commands	
Notes	

11.7.1.12 show qos interface l2-mapping

	show qos interface <type> <port-id> l2-mapping Displays the PCP, DEI to switch priority table.	
Syntax Description	type	Ethernet, port-channel, or mlag-port-channel
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	

Example	`switch (config)# show qos interface ethernet 1/9 l2-mapping`
	
	`PCP,DEI to switch-priority mapping:`
	`----------------------------------------`
	`PCP(DEI)  switch-priority`
	`----------------------------------------`
	`0(0) 0(1)  0`
	`1(0) 1(1)  1`
	`2(0) 2(1)  2`
	`3(0) 3(1)  3`
	`4(0) 4(1)  4`
	`5(0) 5(1)  5`
	`6(0) 6(1)  6`
	`7(0) 7(1)  7`	

Related Commands		

Notes	

11.7.1.13 show qos interface l3-mapping

	show qos interface <type> <port-id> l3-mapping Displays the DSCP to switch priority table.	
Syntax Description	type	Ethernet, port-channel, or mlag-port-channel
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
Example	switch (config)# show qos interface ethernet 1/9 l3-mapping IP PCP,DEI rewrite: enabled DSCP to switch-priority mapping: ---------- DSCP switch-priority ---------- 0 1 2 3 4 5 6 7 0 8 9 10 11 12 13 14 15 1 16 17 18 19 20 21 22 23 2 24 25 26 27 28 29 30 31 3 32 33 34 35 36 37 38 39 4 40 41 42 43 44 45 46 47 5 48 49 50 51 52 53 54 55 6 56 57 58 59 60 61 62 63 7	
Related Commands		
Notes		

11.7.1.14 show qos interface rewrite-mapping

	show qos interface <type> <port-id> rewrite-mapping Displays the rewrite mapping of switch priority to PCP, DEI and DSCP table.	
Syntax Description	type	Ethernet, port-channel, or mlag-port-channel
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
	3.6.8008	Updated example

Example	switch (config)# show qos interface ethernet 1/1 rewrite-mapping PCP,DEI rewrite : disabled IP PCP,DEI rewrite: enable DSCP rewrite : disabled Rewrite mapping (switch-priority to PCP,DEI,DSCP): Egress Interface: Eth1/1 ---------- switch-priority PCP(DEI) DSCP TC ---------- 0 0(0) 0 0 1 1(0) 8 1 2 2(0) 16 2 3 3(0) 24 3 4 4(0) 32 4 5 5(0) 40 5 6 6(0) 48 6 7 7(0) 56 7
Related Commands	
Notes	

11.7.1.15 show qos interface tc-mapping

	show qos interface <type> <port-id> tc-mapping Displays mapping from switch priority to traffic class.	
Syntax Description	type	Ethernet, port-channel, or mlag-port-channel
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
Example	switch (config)# show qos interface ethernet 1/9 tc-mapping Switch Priority to TC mapping: ---------- Switch Priority TC ---------- 0 0 1 1 2 2 3 3 4 4 5 5 6 6 7 7	
Related Commands		
Notes		

11.7.1.16 show qos mapping ingress interface egress interface

	show qos mapping ingress interface <type> <port-id> egress interface <type> <port-id> Displays end to end mapping configuration: ingress to egress.	
Syntax Description	type	Ethernet, port-channel, or mlag-port-channel
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	

	3.8.2000	Updated example

Example

```
switch (config)# show qos mapping ingress interface ethernet 1/8 egress interface ethernet 1/9

Ingress Interface Eth1/8:
  Trust mode             : L2
  Default Switch Priority: 0
  Rewrite PCP,DEI        : disabled
  Rewrite DSCP           : disabled
  Global Rewrite mode    : enable

PCP,DEI and DSCP to switch-priority mapping:
  ------------------------------------------------------------
  PCP,DEI             DSCP                     switch-priority
  ------------------------------------------------------------
  0(0) 0(1)           0 1 2 3 4 5 6 7          0
  1(0) 1(1)           8 9 10 11 12 13 14 15    1
  2(0) 2(1)           16 17 18 19 20 21 22 23  2
  3(0) 3(1)           24 25 26 27 28 29 30 31  3
  4(0) 4(1)           32 33 34 35 36 37 38 39  4
  5(0) 5(1)           40 41 42 43 44 45 46 47  5
  6(0) 6(1)           48 49 50 51 52 53 54 55  6
  7(0) 7(1)           56 57 58 59 60 61 62 63  7

Egress Interface: Eth1/9

----------------------------------------
switch-priority   PCP(DEI)   DSCP     TC
----------------------------------------
0                 0(0)       0        0
1                 1(0)       8        1
2                 2(0)       16       2
3                 3(0)       24       3
4                 4(0)       32       4
5                 5(0)       40       5
6                 6(0)       48       6
7                 7(0)       56       7

r-qa-sw-eth-84 [standalone: master] (config) #
```

Related Commands	
Notes	

11.7.2 QoS Rewrite

11.7.2.1 qos rewrite pcp

	qos rewrite pcp Enables PCP,DEI rewrite on the interface. The no form of the command disables PCP,DEI rewrite on the interface.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.1002	
	3.8.2000	Updated example
Example	`switch (config interface ethernet 1/1) # qos rewrite pcp`	
Related Commands		

Notes	

11.7.2.2 qos rewrite dscp

	qos rewrite dscp Enables DSCP rewrite on the interface. The no form of the command disables DSCP rewrite on the interface.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.1002	
	3.8.2000	Updated example
Example	`switch (config interface ethernet 1/1) # qos rewrite dscp`	
Related Commands		
Notes		

11.7.2.3 qos rewrite map switch-priority pcp dei

	qos rewrite map switch-priority <switch-priority> pcp <pcp> dei <dei> no qos rewrite map switch-priority <switch-priority> pcp <pcp> dei <dei> Configures switch-priority to PCP,DEI mapping on the interface. The no form of the command resets the value to their defaults.	
Syntax Description	switch-priority	Range: 0-7
	pcp	Range: 0-7
	dei	Value: 0
Default	Switch priority to PCP,DEI mapping: 0 → 0,0 1 → 1,0 2 → 2,0 3 → 3,0 4 → 4,0 5 → 5,0 6 → 6,0 7 → 7,0	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.1002	
	3.8.2000	Updated example

Example	`switch (config interface ethernet 1/1) # qos rewrite map switch-priority (0-7) pcp 7 dei 0`
	
	`switch (config interface ethernet 1/14) # no qos rewrite map switch-priority 7 pcp`	

Related Commands		

Notes	

11.7.2.4 qos rewrite map switch-priority dscp

	qos rewrite map switch-priority <switch-priority> dscp <dscp> no qos rewrite map switch-priority <switch-priority> dscp <dscp> Configures switch-priority to DSCP mapping on the interface. The no form of the command resets the value to their defaults.	
Syntax Description	N/A	
Default	Switch priority to DSCP mapping: 0 → 0 1 → 8 2 → 16 3 → 24 4 → 32 5 → 40 6 → 48 7 → 54	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.1002	
Example	`switch (config interface ethernet 1/1) # qos rewrite map switch-priority 5 dscp 40`	
Related Commands		
Notes		

11.7.2.5 qos ip rewrite pcp

	qos ip rewrite pcp [disable | enable | preserve] no qos ip rewrite pcp [disable | enable | preserve] Enables or preserves the rewrite of PCP, DEI of routed packets in egress interface. The no form of the command resets the value to their defaults.	
Syntax Description	disable	No rewrite occurs
	enable	PCP,DEI are rewritten based on the mapping configured on the egress port
	preserve	Ingress interface configuration determines action
Default	Enable	
Configuration Mode	config	
History	3.6.1002	
Example	`switch (config) # qos ip rewrite pcp enable`	
Related Commands		
Notes		

11.7.2.6 show qos ip rewrite

	show qos ip rewrite Displays configuration of the rewrite of PCP, DEI of routed packets in egress interface
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.6000
Example	`switch (config)# show qos ip rewrite` `IP rewrite PCP: enable`
Related Commands	qos ip rewrite pcp
Notes	

11.7.3 Queuing and Scheduling (ETS)

11.7.3.1 bind switch-priority

	bind switch-priority [<priority_1> [<priority_2] .. <priority_n>]] no bind switch-priority [<priority>] Configures binding of switch-priority to traffic class. The no form of the command: • When run in the interface configuration mode: Resets to default the binding of all switch-priorities from all traffic classes • When run in the interface's traffic class: Negates the binding of a specific switch-priority from a specific traffic class
Syntax Description	N/A
Default	Switch priority to traffic class mapping: 0 → 0 1 → 1 2 → 2 3 → 3 4 → 4 5 → 5 6 → 6 7 → 7
Configuration Mode	config interface ethernet config interface ethernet traffic-class config interface port-channel config interface port-channel traffic-class config interface mlag-port-channel config interface mlag-port-channel traffic class
History	3.6.1002
Example	`switch (config 1/1 interface ethernet traffic-class 0) # bind switch-priority 1`
Related Commands	
Notes	Context is egress interface traffic class

11.7.3.2 bandwidth guaranteed

	bandwidth guaranteed [<rate>] no bandwidth guaranteed [<rate>] Configures the minimum bandwidth for outbound traffic. The no form of the command resets this parameter to its default.	
Syntax Description	rate	Rate in GbE Range: 0 - max speed supported
Default	0	
Configuration Mode	config interface ethernet traffic-class config interface port-channel traffic-class config interface mlag-port-channel traffic class	
History	3.6.1002	
Example	`switch (config interface ethernet 1/1 traffic-class 0) # bandwidth guaranteed 0.4G`	
Related Commands		
Notes	• Context is egress interface traffic class • Bandwidth guaranteed rate determines the bandwidth guaranteed by the switch for outbound traffic assigned to this traffic class on this interface • Bandwidth is in granularity of 0.2G	

11.7.3.3 bandwidth shape

	bandwidth shape [<shape>] no bandwidth shape [<shape>] Configures the bandwidth shaper for outbound traffic. The no form of the command resets this parameter to its default.	
Syntax Description	shape	Rate in GbE Range: 0 - max speed supported (in increments of 0.2)
Default	Maximum port rate	
Configuration Mode	config interface ethernet traffic-class config interface port-channel traffic-class config interface mlag-port-channel traffic class	
History	3.6.1002	
	3.9.2000	Updated notes
Example	`switch (config interface ethernet 1/1 traffic-class 7) # bandwidth shape 0.4G`	
Related Commands		
Notes	• Context is egress interface traffic class and/or port • Bandwidth shape rate determines the bandwidth of the shaper for outbound traffic assigned to this traffic class on this interface • Bandwidth is in granularity of 0.2G • Configuring shaping of a LAG group means configuring the same shaper value for each physical port in the LAG • Shaping on a LAG is limited to the LAG member bandwidth	

11.7.3.4 show dcb ets

	show dcb ets [interface {ethernet | mlag-port-channel | port-channel} <if-id>] Displays ETS information.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	
	3.6.5000	Updated example
Example	switch (config)# show dcb ets interface ethernet 1/1 Eth1/1: Interface Bandwidth Shape [Mbps]: N/A Multicast unaware mapping: disabled Flags: S.Mode: Scheduling Mode [Strict/WRR] D: - W: Weight Bw.Sh: Bandwidth Shaper Bw.Gr: Bandwidth Guaranteed ETS per TC: ---------- TC S.Mode W W(%) BW Sh.(Mbps) BW Gr.(Mbps) ---------- 0 WRR 12 12 N/A 0 1 WRR 13 13 N/A 0 2 WRR 12 12 N/A 0 3 WRR 13 13 N/A 0 4 WRR 12 12 N/A 0 5 WRR 13 13 N/A 0 6 WRR 12 12 N/A 0 7 WRR 13 13 N/A 0	
Related Commands		
Notes		

11.7.4 RED & ECN

11.7.4.1 traffic-class congestion-control

	traffic-class <tc> congestion-control [red | ecn | both] [minimum- absolute <min> maximum-absolute <max> | minimum-relative <min> maximum-relative <max>] no traffic-class <tc> congestion-control Enables RED/ECN marking for traffic class queue. The no form of the command disables RED/ECN marking for traffic class queue.	
Syntax Description	tc	Traffic class. Range: 0-7
	red	Enables random early detection for traffic class queue.
	ecn	Enables explicit congestion notification for traffic class queue.
	both	Enables both RED and ECN marking for traffic class queue.
	minimum-absolute	Set minimum-absolute value (in KBs) for marking traffic-class queue.

	maximum-absolute	Set maximum-absolute value (in KBs) for marking traffic-class queue.
	minimum-relative	Set minimum-relative value (in percentage) for marking traffic-class queue.
	maximum-relative	Set maximum-relative value (in percentage) for marking traffic-class queue.
Default	Disabled	
Configuration Mode	config interface ethernet	
History	3.5.1000	
	3.9.1300	Added example

Example

```
switch (config interface ethernet 1/1)# traffic-class 0 congestion-control both minimum-relative 50
maximum-relative 80

2100:
switch (config) # interface ethernet 1/4 traffic-class 3 congestion-control ecn minimum-absolute 12
maximum-absolute ?
12 - 12111 KBs value

3700:
switch (config) # interface ethernet 1/1 traffic-class 4 congestion-control ecn minimum-absolute ?
3 - 30703 KBs value

2700:
switch (config) # interface ethernet 1/1 traffic-class 3 congestion-control ecn minimum-absolute ?
3 - 10863 KBs value

2410:
switch (config) # interface ethernet 1/1 traffic-class 1 congestion-control ecn minimum-absolute ?
3 - 8991


4600: switch (config) # interface ethernet 1/1 traffic-class 4 congestion-control ecn minimum-absolute ?

3 - 54063  KBs value

4700: switch (config) # interface ethernet 1/1 traffic-class 4 congestion-control ecn minimum-absolute ?

3 - 55631 KBs value
```

Related Commands	
Notes	

11.7.4.2 show interfaces ethernet congestion-control

	show interfaces ethernet congestion-control Displays specific interface congestion control information.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.5.1000

Example	switch (config)# show interface ethernet 1/1 congestion-control Interface ethernet: 1/1 ECN marked packets: 0 TC-0 Mode: ECN Threshold mode: absolute Minimum threshold: 0 KB Maximum threshold: 200 KB RED dropped packets: 0 TC-1 Mode: RED Threshold mode: relative Minimum threshold: 0% Maximum threshold: 100% RED dropped packets: 0 TC-2 Mode: none TC-3 Mode: none TC-4 Mode: ECN Threshold mode: relative Minimum threshold: 25% Maximum threshold: 80% RED dropped packets: 0 TC-5 Mode: none TC-6 Mode: both Threshold mode: absolute Minimum threshold: 100 KB Maximum threshold: 200 KB RED dropped packets: 0 TC-7 Mode: none
Related Commands	
Notes	

11.8 Priority Flow Control (PFC)

Priority Flow Control (PFC) provides an enhancement to the existing pause mechanism in Ethernet. The current Ethernet pause option stops all traffic on a link. PFC creates eight separate virtual links on the physical link and allows any of these links to be paused and restarted independently, enabling the network to create a no-drop class of service for an individual virtual link.

PFC offers the following features:

- Provides per-priority enabling or disabling of flow control
- Transmits PFC-PAUSE frames when the receive threshold for a particular traffic class is reached
- Provides the management capability for an administrator to configure the flow control properties on each port of the switch
- Keeps flow control disabled for all priorities on all ports by default
- Allows an administrator to enable or disable flow control per port and per priority level
- Supports flow control only on physical ports, not on logical interfaces such as tunnels or interfaces defined by sharing a physical port in multiple virtual switch contexts
- Uses the configured threshold values to set up the queue buffer spaces accordingly in the data-path

- Provides hardware abstraction layer call-outs for the following:
- Enabling or disabling of flow control on each port for each priority
- Configuring the queue depth for each priority on each port
- Provides trace logs for execution upon error conditions and for any event notifications from the hardware or data-path. These trace logs are a useful aid in troubleshooting.
- Allows the administrator to configure the minimum and maximum threshold values for flow control. These configurations are applied globally on all ports and priorities.

Priority Based Flow Control (PFC) provides an enhancement to the existing pause flow control mechanism as described in 802.1x.

To enable PFC globally,

```
switch (config) # dcb priority-flow-control enable
This action might cause traffic loss while shutting down a port with priority-flow-control mode on
Type 'yes' to confirm enable pfc globally: yes
```

To enable PFC per priority:

1. Enable PFC globally on the switch.

   ```
   switch (config) # dcb priority-flow-control enable
   This action might cause traffic loss while shutting down a port with priority-flow-control mode on
   Type 'yes' to confirm enable pfc globally: yes
   ```

2. Choose the priority you want to enable.

   ```
   switch (config) # dcb priority-flow-control priority 5 enable
   ```

To enable PFC per interface, do the following.

1. Enable PFC globally on the switch.

   ```
   switch (config) # dcb priority-flow-control enable
   ```

2. Choose the priority you want to enable.

   ```
   switch (config) # dcb priority-flow-control 5 enable
   ```

3. Change to Interface mode.

   ```
   switch (config) # interface ethernet 1/1
   switch (config interface ethernet 1/1) #
   ```

4. Enable PFC for the specific interface.

   ```
   switch (config interface ethernet 1/1) # dcb priority-flow-control mode on
   ```

When working with lossless traffic, the receiving side sends a pause frame (Xoff) to the transmitting side before the buffer is filled. When the buffer empties, the receiving side sends an un-pause frame (Xon) to the transmitting side.

11.8.1 Flow Control Threshold Configuration

The user has to set the buffer usage Xoff and Xon thresholds. The thresholds depend on network parameters (bandwidth, link latency, MTU) and the allocated size for the region.

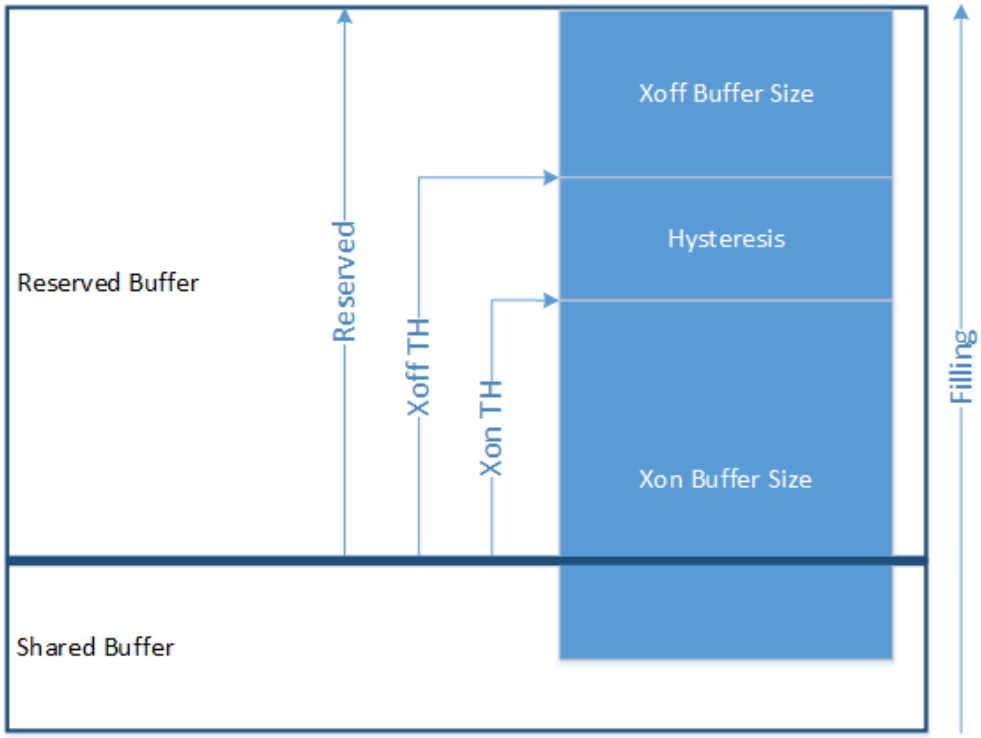

When working with global flow control mode only, a single PG shall be used and Xoff and Xon shall be set on this PG. When working with priority flow control, Xoff and Xon shall be set on each lossless PG.

See the "Shared Buffers" page for more information on flow control.

11.8.2 PFC Watchdog

Lossless networks with PFC enabled provide strong packet delivery guarantees. However, lossless networks introduce a new fault scenario where a queue of an end-port (e.g. the port of a host connected to the network) may not be able to receive any traffic from the network and keeps sending pause frames towards the switch. Since lossless switch paths do not drop packets but decline receiving more packets when their buffers fill up, if the end-port queue is stuck for a long time, the buffers fill up not only for the target switch, but also on all switches with problematic port queues in the traffic forwarding path. This leads to endless PFC pause frames, also called a PFC storm, being observed on all switch ports along the path to the traffic source.

PFC watchdog prevents congestion from spreading in such a case. When switches detect this situation on any TC queue, all the packets in the queue are flushed and new packets destined to the same queue are dropped as well until PFC storming is relieved.

For lossless networks with global flow control configured, we will face the same issue of global pause storm. To resolve this, global-flow-control-watchdog mode is supported.

11.8.3 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community post:

- How to Enable PFC on Switches (Spectrum)

11.8.4 PFC Commands

11.8.4.1 dcb priority-flow-control enable

	dcb priority-flow-control enable [force] disable dcb priority-flow-control [force] no dcb priority-flow-control enable [force] Enables PFC globally on the switch. It is also possible to assign specific interface behavior in dcb priority-flow-control mode. The disable form of the command globally disables PFC on the switch when RoCE mode is set to lossless/semi-lossless. The no form of the command sets global PFC to the default value. See “Default” section below.	
Syntax Description	force	Forces operation
Default	PFC is generally disabled. See “RoCE Parameters” for specific RoCE modes in which the default is enabled	
Configuration Mode	config	
History	3.1.0000	
	3.3.0000	Updated example
	3.8.2100	• Updated example • Added "disable dcb priority-flow-control" command • Changes the function of the no form of the command
	3.9.0500	Updated the description of the "disable" form of the command and added a note

Example	`switch (config)# no roce`
	`switch (config)# no dcb priority-flow-control enable force`
	`switch (config)# show dcb priority-flow-control`
	`PFC: disabled`
	`switch (config)# dcb priority-flow-control enable`
	`This action might cause traffic loss while shutting down a port with priority-flow-control mode on`
	`Type 'yes' to confirm enable pfc globally: yes`
	`switch (config)# show dcb priority-flow-control`
	`PFC: enabled`
	`switch (config)# roce semi-lossless`
	`switch (config)# show dcb priority-flow-control`
	`PFC: enabled`
	`switch (config)# disable dcb priority-flow-control force`
	`switch (config)# show dcb priority-flow-control`
	`PFC: disabled`
	`switch (config)# no dcb priority-flow-control enable force`
	`switch (config)# show dcb priority-flow-control`
	`PFC: enabled`	

Related Commands	show dcb priority-flow-control dcb priority-flow-control mode	

Notes	• This command asks the user to approve traffic loss because some interfaces with DCB mode activated might get shut down. • The disable command is valid only for roce lossless/semi-lossless modes. For explicitly disabling PFC on other scenarios, please set the interface PFC mode to 'off' for all required ports.

11.8.4.2 dcb priority-flow-control priority

	dcb priority-flow-control priority <prio> enable no dcb priority-flow-control priority <prio> enable Enables PFC per priority on the switch. The no form of the command disables PFC per priority on the switch.	
Syntax Description	prio	0-7
Default	PFC is disabled for all priorities.	
Configuration Mode	config	
History	3.1.0000	
	3.9.0500	Added note
Example	`switch (config)# dcb priority-flow-control priority 0 enable`	
Related Commands	show dcb priority-flow-control	
Notes	When RoCE mode is set to lossless/semi-lossless, the no form of the command is not applicable. For explicitly disabling PFC, set the interface PFC mode to 'off' for all required ports.	

11.8.4.3 dcb priority-flow-control mode

	dcb priority-flow-control mode <mode> [force] no dcb priority-flow-control mode [force] Changes PFC mode per interface. The no form of the command disables PFC per interface.	
Syntax Description	force	Configures the PFC admin mode as on or auto with no confirmation needed if the port is admin enabled
	mode	The interface PFC mode. Possible values: • on - enables PFC per interface • off - disables PFC per interface • auto - set PFC mode for the interface to be controlled with traffic pool configuration
Default	auto - PFC mode is established by traffic pool configuration (not a directly configurable mode)	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.3.4500	Added MPO configuration mode
	3.6.6000	Added “force” parameter
	3.6.6102	Added “mode” parameter
	3.6.7100	Updated “mode” parameter description

Example	`switch (config interface ethernet 1/1) # dcb priority-flow-control mode on`
Related Commands	show dcb priority-flow-control
Notes	• For the “force” parameter, the no form of the command disables priority-flow-control without the preceding confirmation prompt • For mode value “auto”, if a lossless traffic pool is configured, PFC is enabled for this port. Otherwise, PFC is disabled.

11.8.4.4 pfc-wd

	pfc-wd no pfc-wd Enables PFC watchdog on interface. The no form of the command disables PFC watchdog on interface.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.6.6000
Example	`switch (config interface ethernet 1/1) # pfc-wd`
Related Commands	show interface pfc-wd
Notes	When a user enables both "flowcontrol receive on" and "pfc-wd" on specific port, global-flow-control-watchdog mode is activated. If only "pfc-wd" is enabled, then the PFC-watchdog mode is activated.

11.8.4.5 show dcb priority-flow-control

	show dcb priority-flow-control [interface <type> <inf>] [detail] Displays DCB priority flow control configuration and status.	
Syntax Description	type	• ethernet • port-channel
	inf	The interface number
	detail	Adds details information to the show output
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	

Example	switch (config) # show dcb priority-flow-control PFC enabled Priority Enabled List : 0 Priority Disabled List : 1 2 3 4 5 6 7 TC Lossless --- ---------- 0 N 1 Y 2 Y 3 N Interface PFC admin PFC oper ------------ -------------- ------------- 1/1 On Enabled 1/2 Disabled Disabled 1/3 Disabled Disabled 1/4 Disabled Disabled ...
Related Commands	
Notes	

11.8.4.6 show dcb priority-flow-control interface mlag-port-channel

	show dcb priority-flow-control interface mlag-port-channel <inf> [detail] Displays DCB priority flow control configuration and status for MPO interfaces.	
Syntax Description	inf	The interface number.
	detail	Adds details information to the show output.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.6.6000	Updated example
Example	switch (config) # show dcb priority-flow-control interface mlag-port-channel 1 detail PFC: disabled Priority Enabled List: Priority Disabled List: 0 1 2 3 4 5 6 7 PFC Port Mpo1 Information: Port Mode : On Operational state : Off No Remote Entry is Present	
Related Commands		
Notes		

11.8.4.7 show interface pfc-wd

	show interface <type> <id> pfc-wd Displays PFC watchdog information.

Syntax Description	type	Interface type: • ethernet • port-channel • mlag-port-channel
	id	Interface ID
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6000	
	3.8.1300	Updated example
Example	switch (config) # show interfaces ethernet 1/1 pfc-wd Interface ethernet 1/1: PFC-WD admin : enable PFC-WD mode : global / per-priority / n/a Traffic Class 0 state: OK Traffic Class 1 state: OK Traffic Class 2 state: OK Traffic Class 3 state: OK Traffic Class 4 state: OK Traffic Class 5 state: OK Traffic Class 6 state: OK Traffic Class 7 state: OK switch (config) #	
Related Commands	pfc-wd	
Notes	When PFC-watchdog mode is activated, display "per-priority" in "PFC-WD mode". While global flow control watchdog activated, display "global". Otherwise, display "n/a".	

11.9 Shared Buffers

All successfully received packets by a switch are stored on internal memory from the time they are received until the time they are transmitted. The packet buffer is fully shared between all physical ports and is hence called a shared buffer. Buffer configuration is applied in order to provide lossless services and to ensure fairness between the ports and priorities.

The buffer mechanism allows defining reserved memory allocation and limiting the usage of memory based on incoming/outgoing ports and priority of the packet. In addition, the buffer can be divided into static pools, each for a specific set of priorities. Buffer configuration mechanism allows fair enforcement from both ingress and egress sides.

The standard configuration mode allows a simple and concise configuration manner by hiding direct buffer access from user, and collecting all the required configuration settings into “traffic pools”. Users that wish to gain full control of entire buffers set can do so by enabling advanced buffer configuration.

11.9.1 Traffic Pool Configuration

The set of configurations which will obtain the optimal shared buffer behavior according to user requirements can be applied by dividing priorities into “traffic pools”. A traffic pool is a logical representation of a traffic profile instance which is supposed to handle all buffer related allocation on the ingress and egress sides to allow fluent flow of the traffic.

Available traffic pool types are as follows:

- Lossy – for standard lossy traffic. This is the default type for all traffic.
- Lossless – for traffic which cannot suffer any loss. Using this type enables a flow control mechanism for the mapped priority as well as setting headroom and Xon/Xoff parameters for the relevant ingress PG buffer.
- Lossy-MC – for layer 2 multicast traffic which requires special care due to stream duplication on the egress side over several ports.

There is no restriction for priority mapping to traffic pools. User can map all priorities to a single traffic pool or create a separate traffic pool for each priority. By default, all memory will be equally divided between all active traffic pools. User can set a memory percentage for a traffic pool out of the entire shared buffer. A state of over-subscription (where sum of percentage is bigger than 100%) is admissible although not advised.

A traffic pool will become functional if at least one priority is mapped to it. Each functional traffic pool will be matched by an iPool, ePool and iPort.PG buffer on each interface. For further detail see section “Advanced Buffer Configuration”.

11.9.2 Lossless Traffic

11.9.2.1 Priority-flow-control

Enabling lossless traffic flow requires relevant switch-priority (see Packet Classification) to be mapped to a traffic pool type “Lossless”. This could be applied through one of the following methods:

- Create a new custom lossless traffic pool, and map the switch-priority to the newly created traffic pool. In this case, PFC configuration is automatic. For example:

```
switch (config) # traffic pool my_pool type lossless
switch (config) # traffic pool my_pool map switch-priority 0
```

- Enabling DCB PFC over the said switch-priority along with enabling DCB PFC globally. This will result in mapping of the priority to the lossless-default traffic pool which is reserved merely for this purpose. In addition it is required to enable DCB PFC for the relevant interfaces as well.

When setting lossless traffic configuration, it is strongly recommended to stick with one of the upper modes rather than a combination of them.

11.9.2.2 Flow Control (Global Pause)

Utilizing global pause mechanism requires “flowcontrol” to be enabled over the desired port and the port's default priority must be set to switch-priority 3 to configure lossless traffic over the port. The configuration steps are described in section “Priority-flow-control”.

To ensure all incoming packets are subjected to the global pause mechanism, the port's trust mode must be set to “port”.

Example:

```
switch (config)# traffic pool my_pool type lossless
switch (config)# traffic pool my_pool map switch-priority 3
switch (config)# interface ethernet 1/1 flowcontrol send on force
switch (config)# interface ethernet 1/1 flowcontrol receive on force
switch (config)# interface ethernet 1/1 qos default switch-priority 3
switch (config)# interface ethernet 1/1 qos trust port
```

11.9.3 Advanced Buffer Configuration

11.9.3.1 Packet Buffering Classification

When a packet arrives to the switch it is classified according to its ingress port, egress port, and layer 2 and layer 3 header fields. The following terms are used to handle packet classification within the switch:

- Port
  - Ingress port (iPort) - the port which the packet is received on
  - Egress port (ePort) - the port which the packet is transmitted on
- Pool
  - Ingress pool (iPool) - the memory pool on which the packet is counted on the ingress side
  - Egress pool (ePool) - the memory pool on which the packet is counted on the egress side
- Priority
  - Switch priority (SP) - internal identifier of the packet priority which is used as a key for several internal switch functions and decisions, including buffering. The SP of the packet is assigned according to a port’s trust level configuration and packet QoS identifiers in the header (PCP, DEI, DSCP).
  - Priority group (PG) - PG is combined of a group of SPs. It is used for grouping packets of several switch priorities into a single ingress buffer space. PG range is from 0-7, while PG 9 is reserved for control traffic.
  - Traffic class (TC) - TC is combined of a group of SPs. It is used for grouping packets of several switch priorities into a single egress queue and buffer space. TC range is from 0-15, while TC 8-15 is reserved for multicast traffic and TC 16 is reserved for control traffic.

Buffer configuration mechanism provides a way to allocate buffer space for specific traffic types by configuring buffers of the following types.

- iPort.PG - traffic which arrives on a specific port and is mapped to a specific PG

- iPort (iPort.pool) – traffic which arrives on a specific port and is counted on a specific iPool. This sums all iPort.PG mapped to the said iPool.
- ePort.TC – traffic which is transmitted on a specific port and mapped to a specific TC
- ePort (ePort.pool) – traffic which is transmitted on a specific port and counted on a specific ePool. It should sum up all ePort.TCs mapped to the said ePool.

Since multicast packets are duplicated among egress ports, to allow consistent packet counting on ingress and egress sides, the following buffers types are used:

- MC.SP – multicast traffic which is classified per specific switch-priority. Counting occurs on egress side prior to packet duplication.
- ePort.mc – multicast traffic which is going to be transmitted on a specific port

11.9.3.2 Buffer Allocation

For the aforementioned classification parameters, a buffering region can be allocated. The buffering region is defined as a set of one of the following: {iPort}, {iPort.pg}, {ePort}, {ePort.TC}, {MC} or {MC.SP}.

For buffer regions, reserved and shared buffering quotas are allocated based on the following configuration parameters:

- Reserved allocation (size) – guaranteed buffering quota for the region which is not shared with other regions
- Shared allocation (shared) – best-effort buffering quota for the region which can be shared with other regions and allocated dynamically. Region usage cannot overflow this quota. Shared allocation can be set using static or dynamic threshold.
- Shared pool – static bound from which the shared space is dynamically allocated

The iPort.PG buffer can be configured to work in one of two modes:

- Lossy – for lossy traffic
- Lossless – for lossless traffic. In this mode, the user must define the flow control thresholds (Xoff, Xon). Reaching Xoff threshold in a PG buffer occupancy will generate “pause” frames to the sender. Reaching Xon threshold ceases “pause” frames transmission. The reserved allocation for this buffer should be at least the value of Xoff to allow sufficient ingress packet buffering for applying Xon/Xoff thresholds.

After initial admittance to headroom buffer—in which its egress port, TC, and ingress PG are defined —a packet is evaluated for eligibility for being stored in the buffer space until it is forwarded.

Buffer eligibility is defined based on the following conditions:

1. If current usage is below allocation thresholds for all four shared:
   - iPort.PG && iPort && ePort.TC && ePort
2. If there is available quota within at least one of the four reserved allocation regions:
   - For lossy traffic: iPort.PG || iPort || ePort.TC || ePort
   - For lossless traffic: ePort.TC || ePort. Ingress check is not performed since all the ingress reserved space is allocated for headroom.

If a packet is not eligible for buffering:

- For lossy traffic: Packet is dropped

- For lossless traffic: Packet stays in headroom on which Xon/Xoff thresholds are applied

11.9.3.3 Pools

Shared buffer space can be statically divided among multiple pools on the ingress side (iPools) and the egress side (ePools). Each buffer is a region that is mapped to a specific pool.

Each pool has the following parameters:

- Size – the total size which is shared among the regions allocated to that pool. The pool's size binds the amount of cumulative shared usage of the regions that are mapped to the pool. The size can be set to infinite value, in which case occupancy of this pool will not be taken into consideration upon admittance of the packet.

  > The pool size does not include the reserved sizes of regions.

- Mode – working mode
  - Static – each region has a static maximum threshold defined in bytes. The user sets the maximum shared quota for this buffer from a specific pool by providing a percentage out of the bounded pool size. If the size is set to infinite, shared quota for mapped buffers gets set in bytes.
  - Dynamic – each region has a dynamic maximal threshold defined as alpha (α) which is the ratio between the current region usage and the pool's free space (equal to the pool usage subtracted from pool size):
    - α accepts the following values 0, 1/128, 1/64, …1/2,1,2,…,64, infinity
    - Buffer acceptance condition is: region_usage < α*free pool space

The port region is counted against the pool to which the PG/TC region of the packet is mapped.

11.9.3.4 Usage Counting

A packet is counted once on the ingress side and on the egress side.

Direction	Traffic Type	Counting Buffers
Ingress		iPort.PG, iPort
Egress	Unicast	ePort.TC, ePort
	Multicast	MC.SP, ePort.mc

11.9.3.5 Control Traffic Buffering

Control packets are buffered in dedicated pools: iPoolCtrl, ePoolCtrl. Furthermore, each port has a set of buffers which are dedicated to control:

- iPort: iPort.iPoolCtrl
- iPort.PG: iPort.pg9
- ePort: ePort.ePoolCtrl
- ePort.TC: iPort.tc16

All control buffers are mapped to control pools and are not configurable.

11.9.3.6 Default Configuration

The default, out-of-box configuration provides the following settings:

Pools:

- iPool0, ePool0 – default pools for all data traffic. Set to dynamic mode with size of the entire shared buffer each.
- iPoolCtrl, ePoolCtrl – dynamic pools dedicated for control with size of 256KB each
- ePool15 – multicast pool with static mode and infinite size

Buffers:

- All buffer configuration (apart from MC.SP) is similar for all ports
- All switch-priorities are mapped to PG0
- Each switch-priority is mapped to a corresponding TC buffer (i-to-i)

Buffer	Reserved	Shared [%/α/Byte]	Pool	Comment
iPort.iPool0	10KB	alpha 8	iPool0 (fixed)	
iPort.iPoolCtrl	0	alpha 8	iPoolCtrl	iPort control buffer
iPort.pg0	0 (20KB headroom)	alpha 8	iPool0	
iPort.pg9	10KB	alpha 8	iPoolCtrl	iPort.pg control buffer
ePort.ePool0	10KB	alpha 8	ePool0 (fixed)	
ePort.ePoolCtrl	0	alpha 8	ePoolCtrl	ePort control buffer
ePort.mc	10KB	90KB	ePool15 (fixed)	Multicast
ePort.tc0-7	1KB	alpha 8	ePool0	
ePort.tc16	1KB	alpha 8	ePoolCtrl	ePort.tc control buffer
MC.SP0-7	0	alpha ¼	ePool0	Global multicast

11.9.3.7 Configuration Example

The following example exhibits how to divide the buffer among traffic priorities in advanced buffer management mode. Assuming that over an out-of-box lossy default configuration is set, the user here configures buffering for lossless traffic classified to switch-priority 1, over Ethernet interfaces 1/1 and 1/5.

The changes on the default configuration are as follows:

- Advanced buffer management is enabled
- Ingress:
  - iPool1 is assigned a size of 13MB
  - Switch-priority is bound to PG1 to allow separate configuration settings
  - PG1 is mapped to selected pool iPool1, classified as lossless and set sufficient headroom (reserved size) of 85KB. Xon/Xoff thresholds are set to 20KB. The shared alpha coefficient is set to 1.

        - iPort.pool1 buffer receives reserved size of 10k and shared coefficient of alpha 1
    - Egress:
        - ePool1 is assigned an infinite size according to recommended lossless traffic settings
        - TC1 (to which switch-priority is mapped by default) is mapped to the selected pool ePool1, and receives reserved size 0 and an infinite shared threshold
        - ePort.mc buffer receives reserved size 0 and an infinite shared threshold
        - ePort.pool1 buffer receives reserved size 0 and an infinite shared threshold
        - MC.SP1 buffer is mapped to egress pool ePool1, and gets reserved size 0 and an infinite shared threshold
    - Finally, priority-flow-control is enabled over switch-priority 1, and over the selected ports

Example:

```
switch (config) # advanced buffer management force
# Pool configuration
switch (config) # pool iPool1 size 13680063 type dynamic
switch (config) # pool ePool1 size inf type static
# Ingress buffer configuration
switch (config) # interface ethernet 1/1 ingress-buffer iPort pool iPool1 reserved 10k shared alpha 1
switch (config) # interface ethernet 1/1 ingress-buffer iPort.pg1 bind switch-priority 1
switch (config) # interface ethernet 1/1 ingress-buffer iPort.pg1 map pool iPool1 type lossless reserved 85k xoff
20k xon 20k shared alpha 1
switch (config) # interface ethernet 1/1 egress-buffer ePort pool ePool1 reserved 0 shared size inf
switch (config) # interface ethernet 1/1 egress-buffer ePort.tc1 map pool ePool1 reserved 0 shared size inf
switch (config) # interface ethernet 1/1 egress-buffer ePort.mc reserved 0 shared size inf
# Egress buffer configuration
switch (config) # interface ethernet 1/5 ingress-buffer iPort pool iPool1 reserved 10k shared alpha 1
switch (config) # interface ethernet 1/5 ingress-buffer iPort.pg1 bind switch-priority 1
switch (config) # interface ethernet 1/5 ingress-buffer iPort.pg1 map pool iPool1 type lossless reserved 85k xoff
20k xon 20k shared alpha 1
switch (config) # interface ethernet 1/5 egress-buffer ePort pool ePool1 reserved 0 shared size inf
switch (config) # interface ethernet 1/5 egress-buffer ePort.tc1 map pool ePool1 reserved 0 shared size inf
switch (config) # interface ethernet 1/5 egress-buffer ePort.mc reserved 0 shared size inf
# MC buffer configuration
switch (config) # pool ePool1 mc-buffer mc.sp1 reserved 0 shared size inf
# PFC configuration
switch (config) # dcb priority-flow-control enable force
switch (config) # dcb priority-flow-control priority 1 enable
switch (config) # interface ethernet 1/1 dcb priority-flow-control mode on
switch (config) # interface ethernet 1/5 dcb priority-flow-control mode on
```

11.9.3.8 Exceptions to Legal Shared Buffer Configuration

The following configurations are permissible in spite of them not being logical since they are useful to the user in specific advanced situations:

- Global scenarios:
    - Traffic pool memory oversubscription (total X%) and Traffic pools with size ‘Auto’ are not allocated.

      In this scenario, two or more traffic pools are configured so the sum of their sizes (specified in the percentage units) is more than 100%. In this case, upon high utilization, traffic “fights” for resources (free pool memory) and can be lost.
    - Switch priority X is mapped to a non-lossless traffic pool, but PFC is enabled on it, or switch priorities X-1,X are mapped to a non-lossless traffic pool, but PFC is enabled on them

      In these scenarios, switch priority X is mapped to a lossy or lossy-MC traffic pool (traffic is not important and traffic loss is allowed), but pause packet generation (PFC) also is enabled over this priority. These cases are allowed if the user expects traffic to be dropped but has enabled PFC to prevent it.
    - Switch priority X is mapped to a lossless traffic pool, but PFC is disabled on it, or Switch priorities X-1,X are mapped to a lossless traffic pool, but PFC is disabled on them

As opposed to the previous scenarios, here the traffic pool is created as lossless, but pause packet generation is disabled. In these cases, the user expects traffic not to have drops, but it can be dropped.

- Per interface scenarios:
  - <if-id> TC X is mapped to more than one traffic pool, or TCs X,X+1 are mapped to more than one traffic pool.
    In these scenarios, traffic class buffers share the same switch priority and are mapped to two different traffic pool. In this cases, with different traffic pool configuration, behavior of traffic is not determined.
  - <if-id> switch priority X is lossless but neither PFC nor FC is not enabled on this interface, or Switch priorities X-1,X are lossless but neither PFC nor FC is enabled on this interface.
    In these scenarios, the user has created a lossless traffic pool and expects that traffic would not be dropped, but pause packet generation (PFC and FC) is disabled on the interface. In these cases, traffic can be dropped.
  - <if-id> has FC enabled, but default priority 0 is not mapped to lossless traffic pool and FC may not be functional.
    In this scenario, global pause packet (FC) generation is enabled on the interface, but default switch priority (traffic arriving to the switch without priority tagging is assigned the default switch priority) is not in lossless traffic pool. In this case, traffic cam be dropped.
  - <if-id> has insufficient headroom allocation to fulfill configuration derived requirements (MTU, speed, cable-length).
    In this scenario, combination of MTU, speed, cable-length, and amount of lossless traffic pools consumes all free headroom memory. In this case, not all required buffers are configured correctly and traffic can be dropped.

11.9.4 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community post:

- Understanding the Alpha Parameter in the Buffer Configuration of Spectrum Switches

11.9.5 Shared Buffer Commands

- Shared Buffer Commands

11.9.6 Shared Buffer Commands

- 11.9.6.1 traffic pool
- 11.9.6.2 type
- 11.9.6.3 map switch-priority
- 11.9.6.4 type map switch-priority

- 11.9.6.5 memory percent
- 11.9.6.6 advanced buffer management
- 11.9.6.7 ingress-buffer
- 11.9.6.8 egress-buffer
- 11.9.6.9 reserved shared size
- 11.9.6.10 pool size type
- 11.9.6.11 pool reserved shared
- 11.9.6.12 map pool type reserved
- 11.9.6.13 bind switch-priority
- 11.9.6.14 description
- 11.9.6.15 pool mc-buffer
- 11.9.6.16 clear buffers pool mc-buffers max-usage
- 11.9.6.17 clear buffers interface ethernet max-usage
- 11.9.6.18 clear buffers interface max-usage
- 11.9.6.19 clear buffers pool max-usage
- 11.9.6.20 clear buffers pool max-usage
- 11.9.6.21 pool description
- 11.9.6.22 cable-length
- 11.9.6.23 show buffers mode
- 11.9.6.24 show buffers status
- 11.9.6.25 show buffers details
- 11.9.6.26 show buffers pools
- 11.9.6.27 show buffers pools mc-buffers
- 11.9.6.28 show traffic pool
- 11.9.6.29 show traffic pool interface ethernet

11.9.6.1 traffic pool

	traffic pool <name> [force] no traffic pool <name> [force] Creates a traffic pool and enters the traffic pool context on prefix mode enabled. The no form of the command deletes a traffic pool.	
Syntax Description	name	String up to 20 characters
	force	Enforces configuration
Default	N/A	
Configuration Mode	config	
History	3.6.5000	
Example	`switch (config)# traffic pool name` `switch (config pool name)#`	
Related Commands		
Notes		

11.9.6.2 type

	type <type> no type <type> Configures the traffic pool type. The no form of the command resets a traffic pool.	
Syntax Description	type	• lossless • lossy • lossy-mc
Default	Lossy	
Configuration Mode	config pool	
History	3.6.5000	
Example	`switch (config pool name)# type lossless`	
Related Commands		
Notes	When using “traffic pool <name> type <type>”, if the traffic pool does not exist then it is created.	

11.9.6.3 map switch-priority

	map switch-priority <list-of-priorities> no map switch-priority <list-of-priorities> Maps switch-priorities to the traffic pool. The no form of the command unmaps switch-priorities.	
Syntax Description	list-of-priorities	Range: 0-7
Default	N/A	
Configuration Mode	config pool	
History	3.6.5000	
Example	`switch (config pool name)# map switch-priority 2 3 1 7`	
Related Commands		
Notes	When using “traffic pool <name> map switch-priority <list-of-priorities>”, if the traffic pool does not exist then it is created.	

11.9.6.4 type map switch-priority

	type {lossless | lossy | lossy-mc} map switch-priority <priority> no type {lossless | lossy | lossy-mc} map switch-priority Configures type of traffic pool and maps switch-priorities to it. The no form of the command unmaps switch-priorities.	
Syntax Description	type	• lossless • lossy • lossy-mc
	priority	Range: 0-7
Default	Type: Lossy	
Configuration Mode	config pool	

History	3.6.5000
Example	`switch (config pool name)# type lossy-mc map switch-priority 2 3 1 7`
Related Commands	
Notes	When using “traffic pool <name> type <type> map switch-priority <priority>”, if the traffic pool does not exist the it is created.

11.9.6.5 memory percent

	memory percent [<percent>] no memory percent [<percent>] Sets traffic pool size in percentage out of entire shared buffer memory. The no form of the command resets this parameter to its default.	
Syntax Description	percent	Range: 0.00-100.00 or “auto”
Default	Auto	
Configuration Mode	config pool	
History	3.6.5000	
Example	`switch (config pool name)# memory percent 50.03`	
Related Commands		
Notes	• Setting “auto” value ensures fair memory division between all traffic pools with “auto” size • Over-subscription of more than 100% is allowed but not recommended, and causes an exception to be displayed in the “Exceptions list” in “show traffic pool” command output. See section “Exceptions to Legal Shared Buffer Configuration” for more details.	

11.9.6.6 advanced buffer management

	advanced buffer management [force] no advanced buffer management [force] Enable the advanced mode shared buffer configuration. The no form of the command disables the advanced mode shared buffer configuration.	
Syntax Description	force	Run command skipping confirmation prompt
Default	Disabled	
Configuration Mode	config	
History	3.6.5000	
	3.6.8008	Updated Note field
Example	`switch (config)# advanced buffer management force` `This will reset all configuration to default. Type ‘yes’ to confirm: yes`	
Related Commands		
Notes	When moving advanced buffer management from disable to enable, buffer/PFC configuration returns all shared buffer configuration to default.	

11.9.6.7 ingress-buffer

	ingress-buffer <buffer-name> no ingress-buffer <buffer-name> Creates and enters the ingress buffer context. The no form of the command deletes an existing buffer.	
Syntax Description	buffer-name	Name of ingress buffer
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.1002	
Example	`switch (config interface ethernet 1/1)# ingress-buffer iPort.pg1` `switch (config interface ethernet 1/1 ingress-buffer iPort.pg1)#`	
Related Commands		
Notes	iPort.pg9 is reserved for control traffic and hence cannot be edited	

11.9.6.8 egress-buffer

	egress-buffer <buffer-name> no egress-buffer <buffer-name> Creates and enters the buffer context. The no form of the command deletes an existing buffer.	
Syntax Description	buffer-name	Name of egress buffer
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.1002	
Example	`switch (config interface ethernet 1/1)# egress-buffer ePort.tc4` `switch (config interface ethernet 1/1 egress-buffer ePort.tc4)#`	
Related Commands		
Notes	ePort.tc16 is reserved for control traffic and hence cannot be edited	

11.9.6.9 reserved shared size

	reserved <value> shared size <size> no reserved <value> Configures the ePort.mc multicast-buffer. The no form of the command resets buffer to default configuration.	
Syntax Description	buffer-name	Name of egress buffer
	value	Amount of reserved memory for buffer in bytes
	shared size	Shared memory in bytes or “infinite”
Default	According to system default OOB configuration	
Configuration Mode	config interface ethernet egress-buffer config interface ethernet ingress-buffer	
History	3.6.5000	

Example	`switch (config 1/1 egress-buffer ePort.mc)# reserved 5k shared alpha 1/128`
Related Commands	
Notes	• ePort.tc16 is reserved for control traffic and hence cannot be edited • It is possible to use “K” and “M” to define shared size

11.9.6.10 pool size type

	pool <pool-name> size <value> type {static | dynamic} no pool <pool-name> size <value> type {static | dynamic} Creates pool. The no form of the command deletes pool.	
Syntax Description	pool-name	Possible values: • ePool0 ... ePool6 • iPool0 ... iPool6
	size	Size of pool in bytes, or “inf” for infinite
Default	According to system default OOB configuration	
Configuration Mode	config	
History	3.6.5000	
Example	`switch (config)# pool iPool2 size 2M type dynamic` `switch (config)# pool iPool2 size static type static`	
Related Commands		
Notes	It is possible to use “K” for kilobytes and “M” for megabytes to define pool size.	

11.9.6.11 pool reserved shared

	pool <pool-name> reserved <reserved> shared <shared units> <shared> no pool <pool-name> Configures the buffer. The no form of the command resets the values to their default.	
Syntax Description	pool-name	Possible values: iPool0-iPool7
	reserved	Amount of reserved memory for the buffer in bytes
	shared units	The amount of shared memory for this buffer Possible values: alpha, max, size • In alpha mode, alpha can have the following values: 0, 1/128, 1/64 ... 1, 2, 4, ... 64, inf • In max mode, the shared size is defined as a percentage of the pool size • In size mode, the shared size is defined in bytes or infinite
Default	According to system default OOB configuration	
Configuration Mode	config interface ethernet egress-buffer config interface ethernet ingress-buffer	
History	3.6.1002	
Example	`switch (config interface ethernet 1/1 ingress-buffer iPort)# pool iPool0 reserved 90K shared alpha 1/8`	
Related Commands		

Notes	

11.9.6.12 map pool type reserved

	map [pool <pool name> type <type> [xoff <xoff-value> xon <xon value>] reserved <reserved size> shared <shared units> <shared size>] Maps iPort.pg buffer to a given pool and sets its reserved and shared sizes. The no form of the command resets buffer to default pool mapping and configuration.	
Syntax Description	pool-name	Possible values: iPool0 ... iPool7
	type	Possible values: lossy, lossless
	reserved size	Amount of reserved memory for the buffer in bytes
	shared units	Possible values: size, alpha, max
	shared size	The amount of shared memory for this buffer • In alpha mode, alpha can have the following values: 0, 1/128, 1/64 ... 1, 2, 4, ... 64, inf • In max mode, the shared size is defined as a percentage of the pool size • In size mode, the shared size is defined in bytes or infinite Shared size depends on type and size of the given pool: • For static pool shared size is in packets • For dynamic pool shared size is in alpha units • For static pool with infinite size only alpha infinite is supported
	xoff	Relevant only on lossless type, Xoff threshold in bytes
	xon	Relevant only on lossless type, Xon threshold in bytes
Default	According to system default OOB configuration	
Configuration Mode	config interface ethernet ingress-buffer	
History	3.6.1002	
	3.6.5000	Updated command syntax
Example	`switch (config interface ethernet 1/9 ingress-buffer iPort.pg5)# map pool iPool6 type lossy reserved 3k shared alpha 2` `switch (config interface ethernet 1/9 ingress-buffer iPort.pg5)# map pool iPool4 type lossless reserved 7k xoff 2k xon 1k shared max 20`	
Related Commands		
Notes	• Xon and Xoff values are in KB and valid only for “lossless” type • It is possible to use “K” and “M” quantifiers to set reserved size	

11.9.6.13 bind switch-priority

	bind switch-priority <list-of-switch-priorities> no bind switch-priority <list-of-switch-priorities> Bind a switch priority (SP) to an ingress buffer. The no form of the command resets this parameter to its default value.	
Syntax Description	list-of-switch-priorities	Possible values: 0-7
Default	According to system default OOB configuration	

Configuration Mode	config interface ethernet ingress-buffer
History	3.6.1002
Example	`switch (config interface ethernet 1/1 ingress-buffer iPort.pg1)# bind switch-priority 0 1`
Related Commands	
Notes	

11.9.6.14 description

	description <description> no description Configures buffer description. The no form of the command deletes buffer description.	
Syntax Description	description	Text string
Default	“”	
Configuration Mode	config interface ethernet egress-buffer config interface ethernet ingress-buffer	
History	3.6.1002	
Example	`switch (config interface ethernet 1/1 ingress-buffer iPort.pg1)# description example`	
Related Commands		
Notes		

11.9.6.15 pool mc-buffer

	pool <pool-name> mc-buffer <buffer> reserved <reserved> shared <shared units> <shared-size> no pool <pool-name> mc-buffer Maps MC-buffer to specified egress pool and sets its reserved and shared sizes. The no form of the command resets the values to their default.	
Syntax Description	mc-buffer	Buffer can have the values mc.sp0, mc.sp1...mc.sp7
	reserved	The amount of shared memory for this buffer
	shared	The amount of shared memory for this buffer • In alpha mode, alpha can have the following values: 0, 1/128, 1/64 ... 1, 2, 4, ... 64, inf • In max mode, the shared size is defined as a percentage of the pool size • In size mode, the shared size is defined in bytes or infinite
Default	N/A	
Configuration Mode	config config interface ethernet egress-buffer	
History	3.6.1002	
	3.6.5000	Added “size” parameter and note
Example	`switch (config)# pool ePool4 mc-buffer mc.sp6 reserved 3k shared size 2K`	

Related Commands	
Notes	• The qualifiers “K” and “M” may be used to set reserved and shared size • The units alpha, max, size is presented to the user according to the pool type “static”, “dynamic” and “size”: • Alpha when pool type is dynamic and size is defined in bytes • Max when pool type is static and size is defined in bytes • Size when pool type is static and size is infinite

11.9.6.16 clear buffers pool mc-buffers max-usage

	clear buffers pool mc-buffers max-usage Clears max-usage statistics for MC.SP (multicast switch priority, mc.sp0 - mc.sp7) shared buffers.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.8.1000
Example	`switch (config)# clear buffers pool mc-buffers max-usage`
Related Commands	
Notes	

11.9.6.17 clear buffers interface ethernet max-usage

	clear buffers interface ethernet <interface name> max-usage Clears max-usage indicator for all buffers of an interface.	
Syntax Description		Name of the interface
Default	N/A	
Configuration Mode	config	
History	3.6.1002	
	3.8.2000	Added the command to the user manual
Example	`switch (config) # clear buffers interface ethernet 1/1 max-usage`	
Related Commands		
Notes		

11.9.6.18 clear buffers interface max-usage

	clear buffers interface max-usage Clears max-usage indicator for all buffers of all interfaces.
Syntax Description	N/A
Default	N/A

Configuration Mode	config	
History	3.6.1002	
	3.8.2000	Added the command to the user manual
Example	`switch (config) # clear buffers interface max-usage`	
Related Commands		
Notes		

11.9.6.19 clear buffers pool max-usage

	clear buffers pool <pool name> max-usage Clears max-usage indicator for a specific pool.	
Syntax Description	pool name	Name of the ingress/egress pool
Default	N/A	
Configuration Mode	config	
History	3.6.1002	
	3.8.2000	Added the command to the user manual
Example	`switch (config) # clear buffers pool iPool2 max-usage`	
Related Commands		
Notes		

11.9.6.20 clear buffers pool max-usage

	clear buffers pool max-usage Clears max-usage indicator for all pools.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config	
History	3.6.1002	
	3.8.2000	Added the command to the user manual
Example	`switch (config) # clear buffers pool max-usage`	
Related Commands		
Notes		

11.9.6.21 pool description

	pool <pool-name> description <description> no pool <pool-name> description Configures the buffer description of a specific pool-name. The no form of the command resets the values to their default.

Syntax Description	pool-name	Possible values: • ePool0 ... ePool7 • iPool0 ... iPool7
	description	String text (20 character max)
Default	“”	
Configuration Mode	config	
History	3.6.1002	
Example	`switch (config)# pool iPool6 description mapped-to-pg3`	
Related Commands		
Notes		

11.9.6.22 cable-length

	cable-length [<meters>] Configures the cable length in meters for the given port.	
Syntax Description	meters	Cable length in meters Range: 5-100,000
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.5000	
Example	`switch (config interface ethernet 1/4)# cable-length 10`	
Related Commands	show interfaces ethernet cable-length	
Notes	• The user may use the quantifier “K” to indicate kilometers (e.g. “cable-length 5K”) • This command is used to calculate the required buffer to sustain the delay caused by the cable length	

11.9.6.23 show buffers mode

	show buffers mode Displays current mode for shared buffers.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.5000
Example	`switch (config)# show buffers mode` `Current mode: user mode`
Related Commands	
Notes	

11.9.6.24 show buffers status

	show buffers status [interfaces ethernet <slot>/<port>] Displays buffer usage status.	
Syntax Description	<slot>/<port>	Ethernet interface
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	
	3.6.5000	Updated example
	3.6.6000	Updated example
	3.8.2000	Updated example

Example

```
------------------------------------------------------------------------------------------
-----------------------------
Interface  Buffer           Pool       Resv      Shared        Usage     MaxUsage  Resv/
Hdrm Usage  Resv/Hdrm MaxUsage
                                       [Byte]    [%/a/Byte]    [Byte]    [Byte]
[Byte]           [Byte]
------------------------------------------------------------------------------------------
-----------------------------
Eth1/1     iPort.iPool0     iPool0     10.0K     alpha 8       0         0         n/
a              n/a
Eth1/1     iPort.iPool1     iPool1     0         alpha 0       0         0         n/
a              n/a
Eth1/1     iPort.iPool2     iPool2     0         alpha 0       0         0         n/
a              n/a
Eth1/1     iPort.iPool3     iPool3     0         alpha 0       0         0         n/
a              n/a
Eth1/1     iPort.iPool4     iPool4     0         alpha 0       0         0         n/
a              n/a
Eth1/1     iPort.iPool5     iPool5     0         alpha 0       0         0         n/
a              n/a
Eth1/1     iPort.iPool6     iPool6     0         alpha 0       0         0         n/
a              n/a
Eth1/1     iPort.iPool7     iPool7     0         alpha 0       0         0         n/
a              n/a
Eth1/1     iPort.iPoolCtrl  iPoolCtrl  0         alpha 8       0         0         n/
a              n/a
Eth1/1     iPort.pg0        iPool0     0         alpha 8       0         0
0              0
Eth1/1     iPort.pg1        iPool0     0         alpha 0       0         0
0              0
Eth1/1     iPort.pg2        iPool0     0         alpha 0       0         0
0              0
Eth1/1     iPort.pg3        iPool0     0         alpha 0       0         0
0              0
Eth1/1     iPort.pg4        iPool0     0         alpha 0       0         0
0              0
Eth1/1     iPort.pg5        iPool0     0         alpha 0       0         0
0              0
Eth1/1     iPort.pg6        iPool0     0         alpha 0       0         0
0              0
Eth1/1     iPort.pg7        iPool0     0         alpha 0       0         0
0              0
Eth1/1     iPort.pg9        iPoolCtrl  10.0K     alpha 8       0         0
0              0
Eth1/1     ePort.ePool0     ePool0     10.0K     alpha 8       0         0         n/
a              n/a
Eth1/1     ePort.ePool1     ePool1     0         alpha 0       0         0         n/
a              n/a
Eth1/1     ePort.ePool2     ePool2     0         alpha 0       0         0         n/
a              n/a
Eth1/1     ePort.ePool3     ePool3     0         alpha 0       0         0         n/
a              n/a
Eth1/1     ePort.ePool4     ePool4     0         alpha 0       0         0         n/
a              n/a
Eth1/1     ePort.ePool5     ePool5     0         alpha 0       0         0         n/
a              n/a
Eth1/1     ePort.ePool6     ePool6     0         alpha 0       0         0         n/
a              n/a
Eth1/1     ePort.ePool7     ePool7     0         alpha 0       0         0         n/
a              n/a
Eth1/1     ePort.mc         ePool15    10.0K     90.0K         0         0         n/
a              n/a
```

Eth1/1	ePort.ePoolCtrl	ePoolCtrl	0	alpha	8	0	0	n/
a	n/a							
Eth1/1	ePort.tc0	ePool0	1.0K	alpha	8	0	0	n/
a	n/a							
Eth1/1	ePort.tc1	ePool0	1.0K	alpha	8	0	0	n/
a	n/a							
Eth1/1	ePort.tc2	ePool0	1.0K	alpha	8	0	0	n/
a	n/a							
Eth1/1	ePort.tc3	ePool0	1.0K	alpha	8	0	0	n/
a	n/a							
Eth1/1	ePort.tc4	ePool0	1.0K	alpha	8	0	0	n/
a	n/a							
Eth1/1	ePort.tc5	ePool0	1.0K	alpha	8	0	0	n/
a	n/a							
Eth1/1	ePort.tc6	ePool0	1.0K	alpha	8	0	0	n/
a	n/a							
Eth1/1	ePort.tc7	ePool0	1.0K	alpha	8	0	0	n/
a	n/a							
Eth1/1	ePort.tc16	ePoolCtrl	1.0K	alpha	8	0	0	n/
a	n/a							

Related Commands	
Notes	Resv/Hdrm Usage/MaxUsage counters specify the usage of reserved buffer set for lossless PG buffers, and of headroom buffer set to fixed 20KB for lossy PG buffers.

11.9.6.25 show buffers details

	show buffers details [<id>] Displays buffer status in details.	
Syntax Description	<slot>/<port>	Ethernet interface
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	
	3.6.5000	Updated Example
	3.7.1000	Updated Example

Example	switch (config)# show buffers details
	
	Flags:
	  Y: Lossy
	  L: Lossless
	  S: Static
	  D: Dynamic
	
	  Shared size is in percent/Bytes for static pool and in alphas for dynamic pool
	
	Interface Eth1/1:
	
	--- 	
	------ 	
	  Buffer Resv Xoff Xon Shared Pool
	Description
	  [Byte] [Byte] [Byte] [%/a/Byte]
	
	--- 	
	------ 	
	  iPort.iPool0(Y) 10.0K - - alpha 8 iPool0(D)
	  iPort.iPool1(Y) 0 - - alpha 0 iPool1(D)
	  iPort.iPool2(Y) 0 - - alpha 0 iPool2(D)
	  iPort.iPool3(Y) 0 - - alpha 0 iPool3(D)
	  iPort.iPool4(Y) 0 - - alpha 0 iPool4(D)
	  iPort.iPool5(Y) 0 - - alpha 0 iPool5(D)
	  iPort.iPool6(Y) 0 - - alpha 0 iPool6(D)
	  iPort.iPool7(Y) 0 - - alpha 0 iPool7(D)
	  iPort.iPoolCtrl(Y) 0 - - alpha 8 iPoolCtrl(D)
	  iPort.pg0(Y) 0 - - alpha 8 iPool0(D)
	  iPort.pg1(Y) 0 - - alpha 0 iPool0(D)
	  iPort.pg2(Y) 0 - - alpha 0 iPool0(D)
	  iPort.pg3(Y) 0 - - alpha 0 iPool0(D)
	  iPort.pg4(Y) 0 - - alpha 0 iPool0(D)
	  iPort.pg5(Y) 0 - - alpha 0 iPool0(D)
	  iPort.pg6(Y) 0 - - alpha 0 iPool0(D)
	  iPort.pg7(Y) 0 - - alpha 0 iPool0(D)
	  iPort.pg9(Y) 10.0K - - alpha 8 iPoolCtrl(D)
	  ePort.ePool0 10.0K - - alpha 8 ePool0(D)
	  ePort.ePool1 0 - - alpha 0 ePool1(D)
	  ePort.ePool2 0 - - alpha 0 ePool2(D)
	  ePort.ePool3 0 - - alpha 0 ePool3(D)
	  ePort.ePool4 0 - - alpha 0 ePool4(D)
	  ePort.ePool5 0 - - alpha 0 ePool5(D)
	  ePort.ePool6 0 - - alpha 0 ePool6(D)
	  ePort.ePool7 0 - - alpha 0 ePool7(D)
	  ePort.mc 10.0K - - 90.0K ePool15(S)
	  ePort.ePoolCtrl 0 - - alpha 8 ePoolCtrl(D)
	  ePort.tc0 1.0K - - alpha 8 ePool0(D)
	  ePort.tc1 1.0K - - alpha 8 ePool0(D)
	  ePort.tc2 1.0K - - alpha 8 ePool0(D)
	  ePort.tc3 1.0K - - alpha 8 ePool0(D)
	  ePort.tc4 1.0K - - alpha 8 ePool0(D)
	  ePort.tc5 1.0K - - alpha 8 ePool0(D)
	  ePort.tc6 1.0K - - alpha 8 ePool0(D)
	  ePort.tc7 1.0K - - alpha 8 ePool0(D)
	  ePort.tc16 1.0K - - alpha 8 ePoolCtrl(D)
	
	switch-priority to Buffers mapping:
	  ------------------------------ 	
	  Switch-priority Buffer
	  ------------------------------ 	
	  0 iPort.pg0
	  1 iPort.pg0
	  2 iPort.pg0
	  3 iPort.pg0
	  4 iPort.pg0
	  5 iPort.pg0
	  6 iPort.pg0
	  7 iPort.pg0

Related Commands	
Notes	

11.9.6.26 show buffers pools

	show buffers pools [pool-name] Displays buffer pool statistics.	
Syntax Description	pool-name	• iPool0-iPool7 • ePool0-ePool7
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	
	3.6.5000	Updated example output
Example	switch (config)# show buffers pools Flags: S - Static, D - Dynamic ---------- Pool Direction Size Usage MaxUsage Description [Byte] [Byte] [Byte] ---------- iPool0 ingress(D) 13.2M 0 576 Lossy-default iPool1 ingress(D) 0 0 0 iPool2 ingress(D) 0 0 0 iPool3 ingress(D) 0 0 0 iPool4 ingress(D) 0 0 0 iPool5 ingress(D) 0 0 0 iPool6 ingress(D) 0 0 0 iPool7 ingress(D) 0 0 0 iPoolCtrl ingress(D) 256.0K 0 0 Control ePool0 egress(D) 13.2M 0 0 Default ePool1 egress(D) 0 0 0 ePool2 egress(D) 0 0 0 ePool3 egress(D) 10.0K 0 0 ePool4 egress(D) 0 0 0 ePool5 egress(D) 0 0 0 ePool6 egress(D) 0 0 0 ePool7 egress(D) 0 0 0 ePool15 egress(S) inf 0 0 Multicast ePoolCtrl egress(D) 256.0K 0 0 Control	
Related Commands		
Notes	When advanced buffer management is disabled, the “Description” field specifies the e/iPool’s relevant traffic pool name.	

11.9.6.27 show buffers pools mc-buffers

	show buffers pools [<pool-name>] mc-buffers Displays global multicast buffers usage status.	
Syntax Description	pool-name	Possible values: ePool0 ... ePool7
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	

Example	switch (config)# show buffers pools ePool4 mc-buffers --- MC-Buffer Pool Resv [Byte] Shared [%/a/Byte] Usage [Byte] MaxUsage [Byte] --- mc.sp0 ePool0 0 alpha 1/4 0 0 mc.sp1 ePool0 0 alpha 1/4 0 0 mc.sp2 ePool0 0 alpha 1/4 0 0 mc.sp3 ePool0 0 alpha 1/4 0 0 mc.sp4 ePool0 0 alpha 1/4 0 0 mc.sp5 ePool0 0 alpha 1/4 0 0 mc.sp6 ePool0 0 alpha 1/4 0 0 mc.sp7 ePool0 0 alpha 1/4 0 0
Related Commands	
Notes	

11.9.6.28 show traffic pool

	show traffic pool [<name>] Displays state and configuration information for a given traffic pool.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.5000
Example	switch (config)# show traffic pool --- --- Traffic Pool [Bytes] Type Memory [%] Switch Priorities Memory actual [Bytes] Usage [KB] Max Usage --- --- lossless-default (RO) lossless auto 0 0 0 lossy-default lossy auto 0, 1, 2, 3, 4, 5, 6, 7 13.7M 0 0 Exception list: N/A
Related Commands	
Notes	• Omission of traffic pool name displays information about all existing traffic pools • The “Exception list” section displays messages to indicate unrecommended configuration. See section “Exceptions to Legal Shared Buffer Configuration” for more details.

11.9.6.29 show traffic pool interface ethernet

	show traffic pool <name> <device/port> interface ethernet <slot>/<port> Displays state and configuration information for the buffers on a given port related to a given traffic pool.	
Syntax Description	<slot>/<port>	Ethernet interface
Default	N/A	

Configuration Mode	Any command mode	
History	3.6.5000	
	3.8.2000	Updated example

Example	<pre>switch (config)# show traffic pool lossy-default interface ethernet 1/1
	
	--- 	
	Switch-priority Ingress buffer Egress buffer
	--- 	
	0 iPort.pg0 ePort.tc0
	1 iPort.pg0 ePort.tc1
	2 iPort.pg0 ePort.tc2
	3 iPort.pg0 ePort.tc3
	4 iPort.pg0 ePort.tc4
	5 iPort.pg0 ePort.tc5
	6 iPort.pg0 ePort.tc6
	7 iPort.pg0 ePort.tc7
	
	--- 	
	Name Memory percent Size (bytes) Usage (bytes) Max Usage
	--- 	
	lossy-default auto 34.9M 0 0
	
	--- 	
	--------- 	
	Ingress buffer Headroom size (bytes) Xon (bytes) Xoff (bytes) Headroom Usage Headroom
	Max Usage
	--- 	
	--------- 	
	iPort.pg0 20.0K N/A N/A 0 0
	
	--- 	
	Direction Pool Usage (bytes) Pool Max Usage (bytes)
	--- 	
	Ingress 0 0
	Egress 0 0
	
	Exception list:
	N/A</pre>	

Related Commands		
Notes	The “Exception list” section displays messages to indicate unrecommended configuration. See section “Exceptions to Legal Shared Buffer Configuration” for more details.	

11.10 Storm Control

Storm control may be enabled on L2 Ethernet ports, LAGs, and MLAGs to monitor inbound traffic to prevent disruptions caused by a broadcast, multicast, or unicast traffic storm on the physical interfaces.

Storm control utilizes a bandwidth-based method to measure traffic where packets exceeding the percentage level specified by the user are dropped.

Users are able to monitor broadcast, unknown unicast, and unregistered multicast traffic while supporting different thresholds for each type or monitor a summary of all the previously mentioned traffic with one threshold.

11.10.1 Storm Control Commands

11.10.1.1 storm-control

	storm-control {<broadcast | unreg-multicast | unknown-unicast> | all} {level <level> | { bits <bits> | bytes <bytes> | packets <packets> [k|m|g]}} [force] no storm-control {<broadcast | unreg-multicast | unknown-unicast> | all} The command enables Storm Control on selected interface. The no form of the command disables Storm Control on selected interface.	
Syntax Description	broadcast | unreg-multicast | unknown-unicast | all	• Each port can support broadcast, unregistered-multicast, unknown-unicast or all configurations • All means one threshold level for all traffic types. It is identical to configuring broadcast, unregistered-multicast and unknown-unicast together.
	level <level> | { bits <bits> | bytes <bytes> | packets <packets> [k|m|g]}	Storm control per traffic type may be configured with different thresholds: • Level – specifies threshold value in percentages from interface speed • Bits – specifies threshold value in bits per second. Must be specified with multiplier k, m, or g. Possible ranges: [1k...999k][1m...999m][1g...200g]. • Bytes – specifies threshold value in bytes per second. May be specified with multiplier k, m, or g. Possible ranges: [128...999][1k...999k][1m...999m][1g...25g]. • Packets – specifies threshold value in packets per second. May be specified with multiplier k, m, or g. Possible ranges: [1...999][1k...999k][1m...999m][1g...2g].
	force	Resolves collisions and applies new configuration
Default	no storm control	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.4006	
	3.6.4110	Updated command syntax, default and configuration mode
	3.6.6000	Added “config interface mlag port channel” configuration mode
	3.7.0000	Added bits/bytes/packets threshold types

Example	`switch (config interface ethernet 1/1) # storm-control broadcast bits 100 m`
	`switch (config interface ethernet 1/1) # storm-control unknown-unicast level 50`
	`switch (config interface ethernet 1/1) # storm-control unreg-multicast packets 900`
	`switch (config interface ethernet 1/1) # storm-control all bytes 1 g`	

Related Commands		
Notes	• The parameter “all” and other configurations are mutually exclusive • Storm control can be configured on a LAG but cannot be configured on LAG members • Storm control cannot be configured on router ports • Storm control cannot be configured on a destination port in a monitoring session • Units are in 10^n. The parameter “k” equals 1000 and not 1024.	

11.10.1.2 show storm-control

	show storm-control [<interface>] The command displays the configuration levels and dropped packets for each traffic type.	
Syntax Description	interface	• Displays configuration and dropped packets on specific interface • If interface is not specified, displays configuration and dropped packets on all interfaces
Default	N/A	
Configuration Mode	config	
History	3.6.4006	
	3.6.4110	Updated example
	3.7.1000	Updated example
Example	switch (config) # show storm-control Interface Eth1/8: Broadcast : 10% Broadcast packets dropped : 0 Unreg-Mcast : N/A Unreg-Mcast packets dropped : N/A Unkn-Ucast : N/A Unkn-Ucast packets dropped : N/A All traffic types : N/A All traffic types packets dropped: N/A	
Related Commands		
Notes		

11.11 Head-of-Queue Lifetime Limit

Head-of-queue (HoQ) lifetime limit (HLL) is a mechanism which allows discarding packets attempting to be transmitted after HLL time from the time that they were ready to be transmitted at the head of the scheduling group.

When HLL_packet2Stall (7 as default) packets encounter HLL drop, the scheduling group enters a stall state. During that state all packets to the sub-group are discarded. The subgroup exits stall state after HLL_time*8.

A counter called HoQ discard packets counts the number of discarded packets due to HLL.

11.11.1 HoQ Commands

11.11.1.1 hll

	hll <max-time> no hll Configures HLL time on this interface. The no form of the command resets HLL time to its default value.

Syntax Description	max-time	Possible values: • <4 | 16 | 32 | 64 | 128 | 256 | 512>ms • <1 | 2>sec • “inf” to disable HLL
Default	512ms	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.5000	
Example	`switch (config interface ethernet 1/10)# hll 512ms`	
Related Commands		
Notes		

11.12 Store-and-Forward

Store-and-Forward is used to describe a functionality where a switch receives a complete packet, stores it, and only then forwards it.

since the switch make forwarding decisions based on the destination address which is at the header of the packet, the switch can make the forwarding decision before receiving the complete packet, this process is called cut-through, the switch forwards part of the packet before receiving the complete packet.

Cut-through allows lower latency and saves buffer space, but if an error occurred in the packet while utilizing cut-through, the packet will be forwarded with an error, alternatively, utilizing store-and-forward allows the switch to drop erroneous packets.

The standard implementation of forwarding mode is for the entire switch; either all ports on a switch are in store-and-forward mode or all ports on a switch are in cut-through mode. NVIDIA implements forwarding mode per egress port, which is a more flexible method and vital in cases where a switch is connected to both a storage device and a compute server among other setups.

11.12.1 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community post:

- Switch Forwarding: "Store and Forward" vs. "Cut-through"

11.12.2 Store-and-Forward Commands

11.12.2.1 switchmode store-and-forward

	switchmode store-and-forward no switchmode store-and-forward disable switchmode store-and-forward Enables global store-and-forward configuration on the switch. The no form of the command removes store-and-forward configuration from the switch and reverts it back to the switch's global configuration. The disable form of the command configures the forwarding mode to cut-through.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.3640	
	3.6.6000	Added "config interface mlag-port-channel" configuration mode
Example	`switch (config)# switchmode store-and-forward`	
Related Commands		
Notes		

12 Ethernet Switching

The following pages provide information on configuring Ethernet (L2) protocols and features.

- Ethernet Interfaces
- Interface Isolation
- Link Aggregation Group (LAG)
- Link Layer Discovery Protocol (LLDP)
- VLANs
- Voice VLAN
- Spanning Tree Protocol
- MAC Address Table
- MLAG
- Link State Tracking
- QinQ
- Access Control List (ACL)
- User Defined Keys
- OpenFlow

12.1 Ethernet Interfaces

Ethernet interfaces have the following physical set of configurable parameters:

- Admin state – enabling or disabling the interface
- Flow control – admin state per direction (send or receive)
- MTU (Maximum Transmission Unit) – 1500-9216 bytes
- Speed – 1/10/40/56/100GbE (depending interface type and system)
- Description – user defined string
- Module-type – the type of the module plugged in the interface

> To use 100GbE QSFP interfaces as 25/10GbE (via QSA adapter), the speed must be manually set with the command “speed 25000” or “speed 10000” respectively under the interface configuration mode.

12.1.1 Breakout Cables

The breakout cable is a unique NVIDIA capability, where a single physical quad-lane QSFPor double density QSFP (QSFP-DD) port can be divided into various types. QSFP can be divided into 2 dual-lane ports or 4 single-lane ports. QSFP-DD can be divided in 2 quard-lane ports or 4 dual-lane ports.It maximizes the flexibility of the end user to use NVIDIA switch with a combination of different interfaces according to the specific requirements of its network. Certain ports cannot be split at all, and there are ports that can be split into 2 ports only (for more information please refer to your Switch Hardware User Manual). Splitting a port changes the notation of that port from x/y to x/y/z with “x/y” indicating the previous notation of the port prior to the split and “z” indicating the

number of the resulting sub-physical port (1,2 or 1,2,3,4). Each sub-physical port is then handled as an individual port. For example, splitting port 10 into 4 lanes gives the following new ports: 1/10/1, 1/10/2, 1/10/3, 1/10/4.

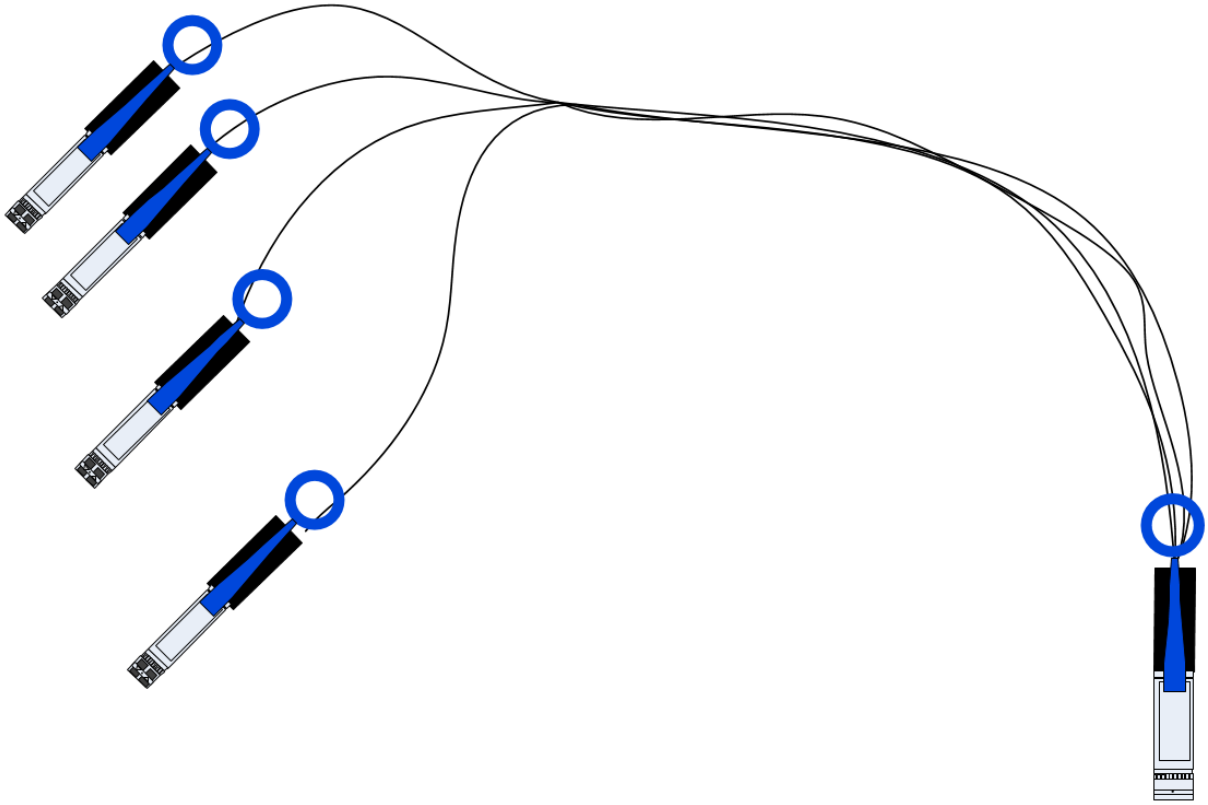

qsfp-split-4 operations result in blocking ports that are additional to the one being split. A set of hardware restrictions determine which of the ports can be split.

Specific ports can be split by using a QSFP 1X4 breakout cable to split one single-lane port into 4 lanes (4 SFP+ connectors). These 4 lanes then go one lane to each of the 4 SFP+ connectors.

> Splitting the interface deletes all configuration on that interface.

When splitting an interface's traffic into 4 data streams (split qsfp into four lanes), one of the other ports on the switch is disabled (unmapped). To see the exact splitting options available per system, refer to the cabling chapter in each specific system's hardware user manual.

To see the exact splitting options available per system, refer to each specific system's hardware user manual (Cabling chapter) located on the company website.

12.1.1.1 Breakout Cables Behavior on SN3800 Switch Systems

SN3800 switch systems currently supports only splitting a port to 2.

12.1.1.2 Changing the Module Type to a Split Mode

To split an interface:

1. Shut down all the ports related to the interface. Run:
   • In case of qsfp-split-2, shut down the current interface only
   • In case of qsfp-split-4, shut down the current interface and the other interface according switch system's specifications.

```
switch (config) # interface ethernet 1/3
switch (config interface ethernet 1/3) # shutdown
switch (config interface ethernet 1/3) # exit
switch (config) # interface ethernet 1/4
switch (config interface ethernet 1/4) # shutdown
```

2. Split
   Split the ports specifying the desired module type to be used (QSFP or QSFP-DD).
   Run:

```
switch (config interface ethernet 1/3) # module-type qsfp-split-4
```

3. The following warning will be displayed:

```
The following interfaces will be unmapped: 1/3 1/4.
Type "YES" when asked to confirm the split.
```

   The <ports> field in the warning refers to the affected ports from splitting port <inf> in the applied command.

> Please beware that in some products splitting a port into a specific type prevents you from accessing the splittable port and an additional one. For example, splitting a port 3 into qsfp-4 on SN2700, makes ports 3 and 4 inaccessible.
>
> This affects the following systems:
>
> • SN2700 – makes one port inaccessible if split port into qsfp-4

12.1.1.3 Unsplitting a Split Port

1. Shut down all of the split ports. Run:

```
switch (config interface ethernet 1/4/4) # shutdown
switch (config interface ethernet 1/4/4) # exit
switch (config) # interface ethernet 1/4/3
switch (config interface ethernet 1/4/3) # shutdown
switch (config interface ethernet 1/4/3) # exit
switch (config) # interface ethernet 1/4/2
switch (config interface ethernet 1/4/2) # shutdown
switch (config interface ethernet 1/4/2) # exit
switch (config) # interface ethernet 1/4/1
switch (config interface ethernet 1/4/1) # shutdown
```

2. From the first member of the split (1/4/1), change the module-type back to QSFP. Run:

```
switch (config interface ethernet 1/4/1) # no module-type
```

> The module-type can be changed only from the first member of the split and not from the interface which has been split.

   The following warning will be displayed:

```
The following interfaces will be unmapped: 1/4/1 1/4/2 1/4/3 1/4/4.
```

3. Type “YES” when prompted with “Type 'YES' to confirm unsplit.”

12.1.2 56GbE Link Speed

NVIDIA offers proprietary speed of 56Gb/s per Ethernet interface.

To achieve 56GbE link speed, run the following on the desired interface:

```
switch (config) # interface ethernet 1/1
switch (config interface ethernet 1/1) # speed 56G
```

For more information about this feature and its potential applications, please refer to the following community post:

- HowTo Configure 56GbE Link on Adapters and Switches

12.1.3 Transceiver Information

NVIDIA Onyx offers the option of viewing the transceiver information of a module or cable connected to a specific interface. The information is a set of read-only parameters burned onto the EEPROM of the transceiver by the manufacture. The parameters include identifier (connector type), cable type, speed and additional inventory attributes.

To display transceiver information of a specific interface, run:

```
switch (config) # show interfaces ethernet 1/20 transceiver
Port 1/20 state
        identifier             : QSFP+
        cable/module type      : Passive copper, unequalized
        ethernet speed and type: 56GigE
        vendor                 : Mellanox
        cable length           : 1m
        part number            : MC2207130-001
        revision               : A3
        serial number          : MT1238VS04936
```

The indicated cable length is rounded up to the nearest natural number.

For more information about this feature and its potential applications, please refer to the following community post:

- HowTo Find Cable Info on Adapters and Switches

12.1.4 High Power Transceivers

NVIDIA switch systems offer high power transceiver (LR4) support in the following ports:

Transceiver			Switch OPN	Supported Ports
Speed	Protocol	Power Consumption [W]		
40GbE	LR4/ER4	3.5	SN2100/SN2410/SN2700	All ports
100GbE		3.5	SN2100/SN2410/SN2700	All ports
100GbE		4.5	SN2100	1, 2, 15, 16
			SN2410	49, 50, 55, 56
			SN2700	1, 2, 31, 32

If a high power transceiver (e.g. LR4) is inserted to a port that does not support it, the link does not go up, and the following warning message is displayed: “Warning: High power transceiver is not supported” when the command “show interfaces ethernet” is run.

12.1.5 Forward Error Correction

Forward Error Correction (FEC) mechanism adds extra data to the transmitted information. The receiving device uses this additional data to verify that the received data contains no errors. If the receiving side discovers errors within the received data it is able to correct some of these errors. The number or errors that can be corrected depends on the FEC algorithm and the amount of redundant data.

100GbE NVIDIA-to-NVIDIA Ethernet connections always enable standard Reed Solomon (RS) FEC on all cables.

If a NVIDIAsystem is connected to a 3rd party system, then FEC is only activated if the 3rd party requests it also.

FEC Modes on All Speeds

Speed	FEC Mode
200GbE	KP4 (enhanced RS FEC)
100/50/25GbE	RS FEC
40/10/1GbE	No FEC

12.1.6 Port Recirculation

The Port Recirculation feature allows the user to configure one of the ports as a recirculation port. When a user configures a physical Ethernet port as a recirculation port, the control of the port will move from the user to the operating system. The interface will no longer be available to the user, but rather be allocated by the operating system for other applications. For instance, on Spectrum-based systems, enabling What-Just-Happened buffer telemetry requires configuring one of the ports

as a recirculation port. In this case, the operating system will use this port to get buffer dropped packets from all the other ports and present them to the user.

12.1.7 Ethernet Interface Commands

- 12.1.7.1 interface ethernet
- 12.1.7.2 boot-delay
- 12.1.7.3 default interface ethernet
- 12.1.7.4 description
- 12.1.7.5 fec-override
- 12.1.7.6 flowcontrol
- 12.1.7.7 ip address dhcp
- 12.1.7.8 load-interval
- 12.1.7.9 module-type
- 12.1.7.10 mtu
- 12.1.7.11 recirculation
- 12.1.7.12 no recirculation port interface ethernet
- 12.1.7.13 shutdown
- 12.1.7.14 speed
- 12.1.7.15 clear counters
- 12.1.7.16 show interfaces counters
- 12.1.7.17 show interfaces counters discard
- 12.1.7.18 show interfaces ethernet
- 12.1.7.19 show interfaces ethernet counters tc
- 12.1.7.20 show interfaces ethernet counters pg
- 12.1.7.21 show interfaces ethernet description
- 12.1.7.22 show interfaces ethernet rates
- 12.1.7.23 show recirculation port
- 12.1.7.24 show interfaces ethernet status
- 12.1.7.25 show interfaces ethernet transceiver
- 12.1.7.26 show interfaces ethernet transceiver brief
- 12.1.7.27 show interfaces ethernet transceiver counters
- 12.1.7.28 show interfaces ethernet transceiver diagnostics
- 12.1.7.29 show interfaces ethernet transceiver raw
- 12.1.7.30 show interfaces status
- 12.1.7.31 disable interface ethernet traffic-class congestion-control
- 12.1.7.32 disable interface port-channel traffic-class congestion-control
- 12.1.7.33 disable interface mlag-port-channel traffic-class congestion-control

12.1.7.1 interface ethernet

	interface ethernet <slot>/<port>[/<subport>][-<slot>/<port>[/<subport>]] Enters the Ethernet interface or Ethernet interface range configuration mode.	
Syntax Description	<slot>/<port>	Ethernet port number
	subport	Ethernet subport number to be used if a port is split

Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.2.1100	Added range support

Example	`switch (config) # interface ethernet 1/1`
	`switch (config) # interface ethernet 1/1) # exit`
	`switch (config) # interface ethernet 1/1-1/10`
	`switch (config) # interface ethernet 1/1-1/10) #`	

Related Commands		
Notes		

12.1.7.2 boot-delay

	boot-delay [<time>] no boot-delay Configures interface boot-delay timer. The no form of the command returns boot-delay time to its default value.	
Syntax Description	time	Boot delay time in seconds Range: 0-600
Default	0 seconds	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.2002	
Example	`switch (config interface ethernet 1/1) # boot-delay 60`	
Related Commands	show interfaces ethernet	
Notes	• This command delays the interface from boot time of the interface • Configuration save and system reboot is required for the configuration to take effect	

12.1.7.3 default interface ethernet

	default interface ethernet <slot/port> Resets a port to its default settings	
Syntax Description	slot/port	Number of Ethernet interface in form of slot/port
Default	N/A	
Configuration Mode	config	
History	3.9.1000	
Example	`switch (config) # default interface ethernet 1/1`	
Related Commands	interfaces ethernet	

Notes	If one of the following configurations exist on the port, the command will be blocked and an informative message will appear. 1. Port is a BGP update source port (when the IP of the port is taken and used as a source IP for BGP routing updates and for TCP connection establishment with neighbor or peer-group). 2. Port is a PIM update source port (when the IP of the port is taken and used as source IP in PIM communications). 3. Port is an IP PIM rp-candidate. 4. Port is an IP PIM bsr-candidate. 5. Port is a member in LAG router port. 6. Port is a member in LAG in NVE mode.

12.1.7.4 description

	description <string> no description Configures an interface description. The no form of the command returns the interface description to its default value.	
Syntax Description	string	40 bytes
Default	""	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.3.4500	Added MPO configuration mode
Example	`switch (config interface ethernet 1/1) # description my-interface`	
Related Commands	show interfaces ethernet	
Notes		

12.1.7.5 fec-override

	fec-override <fec-configuration> [force] no fec-override <fec-configuration> [force] Changes FEC configuration on a specific port or range of ports. The no form of the command resets this parameter to its default value.	
Syntax Description	fec-configuration	• fc-fec – FireCode FEC • no-fec – does not use FEC • rs-fec – Reed Solomon FEC
	force	
Default	Auto-FEC selection	
Configuration Mode	config interface ethernet	
History	3.5.0000	
	3.6.2002	Added force option
	3.7.1000	Updated Example
Example	`switch (config interface ethernet 1/1) # fec-override fc-fec`	
Related Commands	show interfaces ethernet	

Notes	Use this command with caution. There is no limitation in configuring non-standard FEC. It may cause the link to malfunction.

12.1.7.6 flowcontrol

	flowcontrol {receive | send} {off | on} [force] Enables or disables IEEE 802.3x link-level flow control per direction for the specified interface.	
Syntax Description	receive | send	• receive - ingresses direction • send - egresses direction
	off | on	• on - enables IEEE 802.3x link-level flow control for the specified interface on receive or send • off - disables IEEE 802.3x link-level flow control for the specified interface on receive or send
	force	Forces configuration without the need to toggle the interface
Default	receive off; send off	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.3.4500	Added MPO configuration mode
Example	`switch (config interface ethernet 1/1) # flowcontrol receive on`	
Related Commands	show interfaces ethernet	
Notes	To configure global pause please see section "Flowcontrol (Global pause)".	

12.1.7.7 ip address dhcp

	ip address dhcp no ip address dhcp Enables DHCP on this Ethernet interface.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface ethernet set as router interface config interface port-channel set as router interface
History	3.4.2008
Example	`switch (config interface ethernet 1/1) # ip address dhcp`
Related Commands	show interfaces ethernet
Notes	

12.1.7.8 load-interval

	load-interval <time> no load-interval Sets the interface counter interval. The no form of the command resets the interval to its default value.	
Syntax Description	time	In seconds
Default	300 seconds	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.3.0000	
	3.3.4500	Added MPO configuration mode
Example	`switch (config interface ethernet 1/1) # load-interval 30`	
Related Commands	show interfaces ethernet	
Notes	This interval is used for the ingress rate and egress rate counters	

12.1.7.9 module-type

	module-type <type> [force] no module-type <type> [force] Splits the interface into two, four, or eight separate interfaces and specifies the module type intended to be used (QSFP or QSFP-DD). The no form of the command resets the interface to its default configuration (non-split)	
Syntax Description	type	• qsfp-split-2 - port is split into 2 ports using QSFP module, each can run at up to 50GbE • qsfp-split-4 - port is split into 4 ports using QSFP module, each can run at up to 25GbE • qsfp-dd-split-2 - port is split into 2 ports using QSFP-DD module, each can run at up to 200GbE. • qsfp-dd-split-4 - port is split into 4 ports using QSFP-DD module, each can run at up to 100GbE. • qsfp-dd-split-8 - port is split into 8 ports using QSFP-DD module, each can run at up to 50GbE.
	force	Force the split operation without asking for user confirmation.
Default	non-split	
Configuration Mode	config interface ethernet	
History	3.1.1400	
	3.5.0000	Added note
	3.6.3640	Added note
	3.6.4006	Added note
	3.9.0900	• Added QSFP-DD split types • Removed “module-type qsfp” command
	3.10.3100	Added split to 8 for PAM4-based, Spectrum-3 systems.

Example	`switch (config interface ethernet 1/4) # module-type qsfp-split-4`
	`The following interfaces will be unmapped: 1/4 1/1`
	`Type 'YES' to confirm split: YES`

Related Commands	show interfaces ethernet

Notes	• Port cannot be split when storm-control is configured on port
	• Force command don't remove storm-control configuration. Error output:
	% Storm control configuration must be removed from interface Eth1/2
	• After a split port is created or deleted, the forwarding mode for each split port is set according to the global configuration
	• The affected interfaces should be disabled prior to the operation
	• In order to unsplit the interface, use the “no” command.
	• The following speeds are supported on the different Ethernet interface types and depend on the system types and plugged module:
	  • non-split: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 56GbE, 100GbE, 200GbE, 400GbE
	  • qsfp-split-2: 1GbE, 10GbE, 25GbE, 50GbE
	  • qsfp-split-4: 1GbE, 10GbE, 25GbE
	  • qsfp-dd-split-2: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 100GbE, 200GbE,
	  • qsfp-dd-split-4: 1GbE, 10GbE, 25GbE, 50GbE, 100GbE
	  • qsfp-dd-split-8: 1GbE, 10GbE, 25GbE, 50GbE
	• When using split-to-4 and split-to-8, only odd ports can be split. In case split is used, the following even port will be unmapped (e.g., splitting port 1/17 to 4 or to 8 will unmap port 1/18).

12.1.7.10 mtu

	mtu <frame-size> Configures the Maximum Transmission Unit (MTU) frame size for the interface.	
Syntax Description	frame-size	Range: 1500-9216 bytes
Default	9216 bytes	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.3.4500	Added MPO configuration mode
	3.9.2000	Updated default MTU size and added note
Example	`switch (config interface ethernet 1/4) # mtu 9216`	
Related Commands	show interfaces ethernet	
Notes	Switches that perform upgrade to version 3.9.2000, existing interfaces will stay with MTU 1500 (or any other value that was configured). Newly created interfaces (created by split/unsplit operation) will be created with MTU 9216 (the new default). The configured and displayed MTU represents the L3 MTU (being used in IP interfaces). The L2 MTU (being used in physical interfaces) is automatically configured as L3 MTU + 22 Bytes.	

12.1.7.11 recirculation

	recirculation [force] no recirculation Sets the recirculation port. The no form of the command unsets the recirculation port.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config interface ethernet	
History	3.9.0300	
	3.9.1000	Added note
Example	`switch (config interface ethernet 1/1) # recirculation force`	
Related Commands	what-just-happened buffer enable	
Notes	This command reduces by 1 the number of monitor sessions that can be configured. It will fail if the maximum number of monitor sessions are already configured.	

12.1.7.12 no recirculation port interface ethernet

	no recirculation port interface ethernet <port_num> Disables the recirculation port.	
Syntax Description	port_num	Port number
Default	N/A	
Configuration Mode	config	
History	3.9.0300	
Example	`switch (config) # no recirculation port interface ethernet 1/2`	
Related Commands	recirculation show recirculation port	
Notes		

12.1.7.13 shutdown

	shutdown no shutdown Disables the interface. The no form of the command enables the interface.	
Syntax Description	N/A	
Default	Interface is enabled	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.3.4500	Added MPO configuration mode

Example	`switch (config interface ethernet 1/4) # shutdown`
Related Commands	show interfaces ethernet
Notes	

12.1.7.14 speed

	speed {<value> [no-autoneg | speed_value [... speed_value]] | <auto>} [force] no speed Sets the speed of the interface. The no form of the command sets the speed of the interface to its default value.	
Syntax Description	value	The following speeds are available: • 1G or 1000—1GbE • 10G or 10000—10GbE • 25G or 25000—25GbE • 40G or 40000—40GbE • 50G or 50000—50GbE (This speed refers to the speed 50Gx2. See below) • 50Gx1—Port runs at 50Gbps using 1 lane for transmitting (50G PAM4: 1 lane * 50 Gbps) • 50Gx2—Port runs at 50Gbps using 2 lanes for transmitting (50G NRZ: 2 lane * 25 Gbps) • 50GxAuto—Port runs at 50Gbps with auto-select lane count • 56G or 56000—56GbE • 100G or 100000—100GbE (This speed refers to the speed 100Gx4. See below) • 100Gx2—Port runs at 100Gbps using 2 lanes for transmitting (100G PAM4: 2 lanes * 50 Gbps) • 100Gx4—Port runs at 100Gbps using 4 lanes for transmitting (100G NRZ: 4 lanes * 25 Gpbs) • 100GxAuto—Port runs at 100Gbps with auto-select lane count • 200G or 200000—200GbE (This speed refers to the speed 200Gx4. See below) • 200Gx4—Port runs at 200Gbps using 4 lanes for transmitting (200G PAM4: 4 lanes * 50 Gbps) • 400G or 400000—400GbE (This speed refers to the speed 400Gx8. See below) • 400Gx8 - Port runs at 400Gbps using 8 lanes for transmitting (400G PAM4: 8 lanes * 50 Gbps) auto—auto-negotiates link speed (not supported on MPO or LAG interfaces)
	no-autoneg	Disallows auto negotiation link speed on the interface (not supported on MPO or LAG interfaces)
	force	Forces speed change configuration
Default	Depends on the port module type (see the “Notes” section below)	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.5.0000	Added 25GbE, 50GbE, and 100GbE speeds and updated notes
	3.6.6000	Added no-autoneg parameter
	3.9.0600	Removed no-autoneg parameter

	3.9.1000	Updated notesAdded speed with lane configuration
	3.9.2000	Added no-autoneg parameter
Example	switch (config interface ethernet 1/1) # speed 40G switch (config interface ethernet 1/2) # speed 40G no-autoneg switch (config interface ethernet 1/3) # speed 25G no-autoneg force	
Related Commands	show interfaces ethernet	
Notes	• The default speed of an interface depends on its speed capabilities. • It is not possible to set the speed on a LAG or MPO interface • Not all interfaces support all speed options • It is not possible to set “auto” speed along with specific speeds • A port with more than one speed advertised or a port configured to “auto” speed cannot be added to LAG • To change the speed of a LAG interface: a. Remove Ethernet ports from LAG. b. Shutdown ports. c. Reconfigure port speed. d. Re-enable ports. e. Re-add ports to LAG interface. • Speed configuration with lane count affects the Spectrum-2 and Spectrum-3 systems only.	

12.1.7.15 clear counters

	clear counters Clears the interface counters.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.3.4500	Added MPO configuration mode
Example	switch (config interface ethernet 1/1) # clear counters	
Related Commands	show interfaces ethernet	
Notes	This command also clears NVE counters	

12.1.7.16 show interfaces counters

	show interfaces <type> <id> counters Displays the extended counters for the interface.	
Syntax Description	id	Interface number: <slot>/<port> **or** <slot>/<port>-<slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	

	3.6.1002	Added “error packets” counter to Tx
	3.6.4006	Added extended output for storm-control
	3.6.5000	Added hoq discard packets counter
	3.9.0500	Removed Priority option
	3.9.1000	Added ability to use a range of ports and added "ECN marked packets" counter
	3.8.1300	Added note
Example	switch (config) # show interfaces ethernet 1/1-1/2 counters Eth1/1: Rx: 0 packets 0 unicast packets 0 multicast packets 0 broadcast packets 0 bytes 0 packets of 64 bytes 0 packets of 65-127 bytes 0 packets of 128-255 bytes 0 packets of 256-511 bytes 0 packets of 512-1023 bytes 0 packets of 1024-1518 bytes 0 packets Jumbo 0 discard packets 0 error packets 0 fcs errors 0 undersize packets 0 oversize packets 0 pause packets 0 unknown control opcode 0 symbol errors 0 discard packets by storm control Tx: 0 packets 0 unicast packets 0 multicast packets 0 broadcast packets 0 bytes 0 discard packets 0 error packets 0 hoq discard packets 0 pause packets 0 pause duration 0 ECN marked packets Eth1/2: . . .	
Related Commands		
Notes	• NVIDIA-Spectrum-based systems display queue depth for TC0-TC7 • As of version 3.9.1000, the "id" attribute is optional. If nothing is selected, information for all ports will be displayed • Discard Packets counter refers to discards due to insufficient buffer in both RX and TX	

12.1.7.17 show interfaces counters discard

	show interfaces <type> <id> counters discard Displays discarded counters of the interface.	
Syntax Description	id	Interface number: <slot>/<port> **or** <slot>/<port>-<slot>/<port>
Default	N/A	

Configuration Mode	Any command mode	
History	3.6.6102	
	3.9.1000	Made "id" attribute optional
Example	switch (config) # show interfaces ethernet 1/24 counters discard Interface Eth1/24: Rx: 0 discard packets 0 error packets 0 fcs errors 0 undersize packets 0 oversize packets 0 pause packets 0 unknown control opcode 0 symbol errors 0 discard packets by storm control 0 general discard packets 0 policy discard packets 0 invalid tag packets 0 discard packets by vlan filter Tx: 1154059970 discard packets 0 error packets 0 hoq discard packets 0 oversize packets 0 policy discard packets 0 SLL discard packets 11500 no buffer discard mc packets 0 discard packets by vlan filter 0 discard packets by stp filter 0 discard packets by loopback filter	
Related Commands		
Notes	• Discard Packets counter refers to discards due to insufficient buffer in both RX and TX. • The "id" attribute is optional. If nothing is selected, information for all ports will be displayed	

12.1.7.18 show interfaces ethernet

	show interfaces ethernet <inf> [cable-length | capabilities | congestion-control | counters | description | link-diagnostics | pfc-wd | signal-degrade | status | switchport | transceiver] Displays the configuration and status for the interface.	
Syntax Description	inf	Interface number: <slot>/<port> **or** <slot>/<port>-<slot>/<port>
	cable-length	Display cable-length of specific interfaces
	capabilities	Display specific interfaces capabilities information
	congestion-control	Display specific interface congestion control information
	counters	Display specific interfaces counters
	description	Display specific interfaces description information
	link-diagnostics	Display interfaces link diagnostics information
	pfc-wd	Display pfc-wd information
	signal-degrade	Display interfaces signal degrade information
	status	Display specific interfaces status information

	switchport	Display specific interface VLAN-membership information
	transceiver	Display detailed cable info for this port
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.6.1002	Added “error packets” counter to Tx, “Last change in operational status”, and “Isolation group” to output
	3.6.2002	Added “boot delay” parameters to output
	3.6.3640	Added support for “forwarding mode”
	3.6.4110	Updated Example with “Forwarding mode”
	3.6.5000	Added telemetry to output
	3.6.6000	Added output line for “auto-negotiation”
	3.6.8008	Updated example
	3.6.8100	Updated example
	3.7.1100	Updated example and notes
	3.9.1000	Added ability to use a range of ports and updated example

Example	switch (config) # show interfaces ethernet 1/1
	
	Eth1/1:
	
	  Admin state : Disabled
	  Operational state : Down
	  Last change in operational status: Never
	  Boot delay time : 0 sec
	  Description : N/A
	  Mac address : 98:03:9b:94:d9:a0
	  MTU : 1500 bytes (Maximum packet size 1522 bytes)
	  Fec : auto
	  Operational Fec : no-fec
	  Flow-control : receive off send off
	  Supported speeds : 1G 10G 25G 40G 50Gx1 50Gx2 100Gx2 100Gx4 200Gx4 400Gx8
	  Advertised speeds : 100Gx4
	  Actual speed : Unknown
	  Auto-negotiation : Enabled
	  Width reduction mode : Unknown
	  Switchport mode : access
	  MAC learning mode : Enabled
	  Forwarding mode : inherited cut-through
	
	  FCS Ingress : Enabled CRC check
	  FCS Egress : Enabled CRC recalculate
	  FCS Timestamping : Enabled
	
	  Telemetry sampling: Disabled TCs: N/A
	  Telemetry threshold: Disabled TCs: N/A
	  Telemetry threshold level: N/A
	
	  Last clearing of "show interface" counters: Never
	  60 seconds ingress rate : 0 bits/sec, 0 bytes/sec, 0 packets/sec
	  60 seconds egress rate : 0 bits/sec, 0 bytes/sec, 0 packets/sec
	
	  Rx:
	  0 packets
	  0 unicast packets
	  0 multicast packets
	  0 broadcast packets
	  0 bytes
	  0 discard packets
	  0 error packets
	  0 fcs errors
	  0 undersize packets
	  0 oversize packets
	  0 pause packets
	  0 unknown control opcode
	  0 symbol errors
	  0 discard packets by storm control
	
	  Tx:
	  0 packets
	  0 unicast packets
	  0 multicast packets
	  0 broadcast packets
	  0 bytes
	  0 discard packets
	  0 error packets
	  0 hoq discard packets

Related Commands	
Notes	• If a high power transceiver (e.g. LR4) is inserted to a port that does not support it, the link does not go up, and the following warning message is displayed: “Warning: High power transceiver is not supported” when running the command “show interfaces ethernet” is run. For more information, please refer to “High Power Transceivers”. • “Operational Fec” appears as N/A while port is DOWN, and as no-fec/fc-fec/rs-fec while port is UP • As of version 3.9.1000, the "inf" attribute is optional. If nothing is selected, information for all ports will be displayed • The speed with lane count information refers to the Spectrum-2 and Spectrum-3 systems only.

12.1.7.19 show interfaces ethernet counters tc

	show interfaces ethernet [<slot/port> | <slot/port>-<slot/port>] counters tc <priority> Displays traffic class counters for the specified interface and priority.	
Syntax Description	slot/port	Number of Ethernet interface in form of slot/port
	priority	Valid priority values: 0-7 or all
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
	3.9.1000	Added ability to use a range of ports

Example	`switch (config) # show interfaces ethernet 1/1-1/2 counters tc 3`
	
	`Eth1/1:`
	`TC 3`
	`  0  packets`
	`  0  bytes`
	`  0  queue depth`
	`  0  unicast no buffer discard`
	`  0  WRED discard`
	
	`Eth1/2:`
	`TC 3`
	`  0  packets`
	`  0  bytes`
	`  0  queue depth`
	`  0  unicast no buffer discard`
	`  0  WRED discard`	

Related Commands		
Notes	As of version 3.9.1000, the "slot/port" attribute is optional. If nothing is selected, information for all ports will be displayed	

12.1.7.20 show interfaces ethernet counters pg

	show interfaces ethernet [<slot/port> | <slot/port>-<slot/port>] counters pg <priority> Displays priority group counters for the specified interface and priority.	
Syntax Description	slot/port	Number of Ethernet interface in form of slot/port
	priority	Valid priority values: 0-7 or all
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
	3.9.1000	Added ability to use a range of ports

Example	switch (config) # show interfaces ethernet 1/1-1/2 counters pg 3
	
	Eth1/1:
	  PG 0:
	  0 packets
	  0 bytes
	  0 queue depth
	  0 no buffer discard
	  0 shared buffer discard
	
	Eth1/2:
	  PG 0:
	  0 packets
	  0 bytes
	  0 queue depth
	  0 no buffer discard
	  0 shared buffer discard

Related Commands	
Notes	As of version 3.9.1000, the "slot/port" attribute is optional. If nothing is selected, information for all ports will be displayed

12.1.7.21 show interfaces ethernet description

	show interfaces ethernet [<inf>] description Displays the admin status and protocol status for the specified interface.	
Syntax Description	inf	Interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.4.1100	Updated example
	3.8.2000	Updated example
	3.9.1000	Updated example

Example	switch (config) # show interfaces ethernet description
	
	----- 	
	Interface Admin Operational Switchport Speed
	Description
	  state state mode
	----- 	
	Eth1/20 Enabled Up hybrid 10G -
	Eth1/21 Enabled Up hybrid 100Gx4 (auto) -
	Eth1/22 Enabled Up hybrid 100Gx4 (auto) -
	
	switch (config) # show interfaces ethernet 1/20 description
	
	----- 	
	Interface Admin Operational Switchport Speed
	Description
	  state state mode
	----- 	
	Eth1/20 Enabled Up hybrid 50Gx2 -	

Related Commands		
Notes	The speed with lane count information refers to the Spectrum-2 and Spectrum-3 systems only.	

12.1.7.22 show interfaces ethernet rates

	show interfaces ethernet rates [<transfer-rate-unit>] Displays the current transfer rate of the interface.	
Syntax Description	transfer-rate-unit	• bytes - displays interface transfer rates in B/s dynamically (while converting to K/M/G if needed) • KB - displays interface transfer rate in Kb/s • MB - displays interface transfer rate in Mb/s • GB - displays interface transfer rate in Gb/s • bits - displays interface transfer rates in b/s dynamically (while converting to K/M/G if needed) • Kb - displays interface transfer rate in Kb/s • Mb - displays interface transfer rate in Mb/s • Gb - displays interface transfer rate in Gb/s • If no parameter is entered, transfer rate is displayed in bits
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.2002	
	3.7.0000	Added new rates to “transfer-rate-unit”
Example	switch (config) # show interfaces ethernet rates KB Port egress ingress avg rate (KB/s) pkts/sec avg rate (KB/s) pkts/sec --------- ---------- -------- ---------- -------- Eth1/1 0 0 0.032 1 Eth1/2 0 0 0.032 1 Eth1/3 0 0 0 0 ...	
Related Commands		
Notes		

12.1.7.23 show recirculation port

	show recirculation port Shows recirculation port status and information.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.9.0300
Example	switch (config) # show recirculation port
Related Commands	recirculation
Notes	

12.1.7.24 show interfaces ethernet status

	show interfaces ethernet [<inf>] status Displays the status, speed and negotiation mode of the specified interface.

Syntax Description	inf	Interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.4.1100	Updated example
	3.9.1000	Updated example
Example	switch (config) # show interfaces ethernet status Port Operational state Speed Negotiation ---- ----------------- ----- ----------- Eth1/58 Down 40 Gbps No-Negotiation Eth1/59 Up 100Gx4 (auto) Auto Eth1/60 Down (Suspend) 40 Gbps No-Negotiation	
Related Commands		
Notes	The speed with lane count information refers to the Spectrum-2 and Spectrum-3 systems only.	

12.1.7.25 show interfaces ethernet transceiver

	show interfaces ethernet [<inf>] transceiver Displays transceiver information.	
Syntax Description	inf	Interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example	switch (config) # show interfaces ethernet status Port Operational state Speed Negotiation ---- ----------------- ----- ----------- Eth1/58 Down 40 Gbps No-Negotiation Eth1/59 Up 40 Gbps No-Negotiation Eth1/60 Down (Suspend) 40 Gbps No-Negotiation	
Related Commands	switch (config) # show interfaces ethernet 1/1 transceiver Port 1/1 state identifier : QSFP+ cable/module type : Optical cable/module ethernet speed and type: 40GBASE - SR4 vendor : Mellanox cable_length : 50 m part number : MC2210411-SR4 revision : A1 serial number : TT1151-00006	
Notes	• For a full list of the supported cables and transceivers, please refer to the LinkX™ Cables and Transceivers webpage • If a high power transceiver (e.g. LR4) is used, it will be indicated in the field “cable/module type”	

12.1.7.26 show interfaces ethernet transceiver brief

	show interfaces ethernet [<inf>] transceiver brief Display brief transceiver information.	
Syntax Description	inf	Interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6102	
Example	switch (config) # show interfaces ethernet 1/1 transceiver brief show interfaces ethernet transceiver brief ---------- ---------- ---------- ---------- ---------- Interface Identifier Vendor PN SN Rev ---------- ---------- ---------- ---------- ---------- Eth1/1 Eth1/2 QSFP+ Mellanox MCP1600-E00A MT1710VS06916 A3 Eth1/3 QSFP+ Mellanox MCP1600-E00A MT1710VS06929 A3 Eth1/4 QSFP+ Mellanox MCP1600-E00A MT1710VS06953 A3 Eth1/5 QSFP+ Mellanox MCP1600-E00A MT1710VS06923 A3	
Related Commands		
Notes	• For a full list of the supported cables and transceivers, please refer to the LinkX™ Cables and Transceivers webpage • If a high power transceiver (e.g. LR4) is used, it will be indicated in the field “cable/module type”	

12.1.7.27 show interfaces ethernet transceiver counters

	show interfaces ethernet [<inf>] transceiver counters Displays PHY counters related to operational FEC mode and actual number of lanes in the current port.	
Syntax Description	inf	Interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	
	3.10.3000	Updated command description, notes, and example

Example	switch (config) # show interfaces ethernet 1/1 transceiver counters Rx phy received bits 398503339622400 phy symbol errors 0 phy corrected bits 0 Time since last clear 0 Edpl/bip errors lane0 0 Edpl/bip errors lane1 0 Edpl/bip errors lane2 0 Edpl/bip errors lane3 0 raw_ber_magnitude 9 raw_ber_coef 1 effective_ber_magnitude 255 effective_ber_coef 15
Related Commands	
Notes	• The counter “phy received bits” provides information on the total amount of traffic received and can be used to estimate the ratio of error traffic • The counter “phy symbol errors” provides information on the error traffic that was not corrected because the FEC algorithm could not do it or because FEC was not active on this interface • The counter “phy corrected bits” provides the number of corrected bits by the active FEC mode (RS/FC) • The counter "Time since last clear" provides the time passed since the last counters clear event in msec • The counter "Edpl/bip errors lane*" provides the NO FEC corrected counter value on appropriate lane • The counter "FC corrected blocks lane*" provides the FC FEC corrected counter value on appropriate lane • The counter "RS corrected symbols lane*" provides the RS FEC corrected counter value on appropriate lane • The counter "raw_ber_magnitude" provides the BER magnitude value for total number of errors. Used for BER calculation - Raw_BER = raw_ber_coef * 10^(-raw_ber_magnitude) • The counter "raw_ber_coef" provides the BER coefficient value for total number of errors • The counter "effective_ber_magnitude" provides the BER magnitude value for effective number of errors. Used for BER calculation - Effective_BER = effective_ber_coef * 10^(-effective_ber_magnitude) • The counter "effective_ber_coef" provides the BER coefficient value for effective number of errors

12.1.7.28 show interfaces ethernet transceiver diagnostics

	show interfaces ethernet [<inf>] transceiver diagnostics Displays cable channel monitoring and diagnostics info for this interface. Tx and Rx power are reported in mW and dBm units.	
Syntax Description	inf	Interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.2002	
	3.6.4006	Updated Example to report Tx and Rx power in mW and dBm units
	3.6.6000	Updated Example

Example	switch (config) # show interfaces ethernet 1/5 transceiver diagnostics
	
	Port 1/5 transceiver diagnostic data:
	  Temperature (-127C to +127C):
	  Temperature : 26 C
	  Hi Temp Alarm Thresh : 80 C
	  Low Temp Alarm Thresh: -10 C
	  Temperature Alarm : None
	
	  Voltage (0 to 6.5535 V):
	  Voltage : 3.28980 V
	  Hi Volt Alarm Thresh : 3.50000 V
	  Low Volt Alarm Thresh: 3.10000 V
	  Voltage Alarm : None
	
	  Tx Bias Current (0 to 131 mA):
	  Ch1 Tx Current : 6.60000 mA
	  Ch2 Tx Current : 6.60000 mA
	  Ch3 Tx Current : 6.60000 mA
	  Ch4 Tx Current : 6.60000 mA
	  Hi Tx Crnt Alarm Thresh : 8.50000 mA
	  Low Tx Crnt Alarm Thresh: 5.49200 mA
	  Ch1 Tx Current Alarm : None
	  Ch2 Tx Current Alarm : None
	  Ch3 Tx Current Alarm : None
	  Ch4 Tx Current Alarm : None
	
	  Tx Power (0 mW to 6.5535 mW / 8.1647 dBm):
	  Ch1 Tx Power : 1.01420 mW / 0.06124 dBm
	  Ch2 Tx Power : 0.96740 mW / -0.14394 dBm
	  Ch3 Tx Power : 0.96730 mW / -0.14439 dBm
	  Ch4 Tx Power : 0.96050 mW / -0.17503 dBm
	  Hi Tx Power Alarm Thresh : 3.46730 mW / 5.39991 dBm
	  Low Tx Power Alarm Thresh: 0.07240 mW / -11.40261 dBm
	  Ch1 Tx Power Alarm : None
	  Ch2 Tx Power Alarm : None
	  Ch3 Tx Power Alarm : None
	  Ch4 Tx Power Alarm : None
	
	  Rx Power (0 mW to 6.5535 mW / 8.1647 dBm):
	  Ch1 Rx Power : 0.99160 mW / -0.03663 dBm
	  Ch2 Rx Power : 1.06080 mW / 0.25633 dBm
	  Ch3 Rx Power : 1.09810 mW / 0.40642 dBm
	  Ch4 Rx Power : 0.97500 mW / -0.10995 dBm
	  Hi Rx Power Alarm Thresh : 3.46730 mW / 5.39991 dBm
	  Low Rx Power Alarm Thresh: 0.04670 mW / -13.30683 dBm
	  Ch1 Rx Power Alarm : None
	  Ch2 Rx Power Alarm : None
	  Ch3 Rx Power Alarm : None
	  Ch4 Rx Power Alarm : None
	
	  Vendor Date Code (dd-mm-yyyy): 07-11-2016

Related Commands	
Notes	This example is for a QSFP transceiver

12.1.7.29 show interfaces ethernet transceiver raw

	show interfaces ethernet [<inf>] transceiver raw Displays cable info for this interface.	
Syntax Description	inf	Interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	
Example		

```
switch (config) # show interfaces ethernet 1/7 transceiver raw

Port 1/7 raw transceiver data:

I2C Address 0x50, Page 0, 0:255:
  0000  0d 02 06 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  0010  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  0020  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  0030  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  0040  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  0050  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  0060  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  0070  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  0080  0d 00 23 08 00 00 00 00 00 00 00 05 8d 00 00 00  ..#.............
  0090  00 00 01 a0 4d 65 6c 6c 61 6e 6f 78 20 20 20 20  ....Mellanox
  00a0  20 20 20 20 0f 00 02 c9 4d 43 32 32 30 37 31 33      ....MC220713
  00b0  30 2d 30 30 41 20 20 20 41 33 02 03 05 00 46 66  0-00A   A3....Ff
  00c0  00 00 00 00 4d 54 31 32 32 37 56 53 30 30 36 34  ....MT1227VS0064
  00d0  32 20 20 20 31 32 30 37 30 38 20 20 00 00 00 e4  2   120708  ....
  00e0  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  00f0  00 00 00 00 00 00 00 00 00 00 02 00 00 30 00 00
I2C Address 0x50, Pages 1, 128:255:
  0080  0d 02 06 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  0090  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  00a0  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  00b0  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  00c0  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  00d0  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  00e0  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
  00f0  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```

Related Commands	
Notes	

12.1.7.30 show interfaces status

	show interfaces status Displays the configuration and status for the interface.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4006	
	3.9.0300	Updated example—added MTU column
	3.9.1000	Updated example
Example		

```
switch (config) # show interfaces status
-------------------------------------------------------------------------------------------------
Port             Operational state     Admin             Speed              MTU     Description

-------------------------------------------------------------------------------------------------
mgmt0            Up                    Enabled           1000Mb/s (auto)    1500      -
Eth1/1           Down                  Disabled          Unknown            1500      -
Eth1/2           Up                    Enabled           40G                1500      -
Eth1/3           Up                    Enabled           40G                1500      -
Eth1/4           Up                    Enabled           40G                1500      -
Eth1/5           Up                    Enabled           40G                1500      -
Eth1/6           Up                    Enabled           10G                1500      -
Eth1/7           Up                    Enabled           10G                1500      -
Eth1/8           Up                    Enabled           10G                1500      -
Eth1/9           Up                    Enabled           10G                1500      -
Eth1/10          Up                    Enabled           100Gx4             1500      -
Eth1/11          Up                    Enabled           100Gx4             1500      -
Eth1/12          Up                    Enabled           100Gx4             1500      -
Eth1/13          Up                    Enabled           100Gx4             1500      -
Eth1/14          Down                  Disabled          Unknown            1500      -
Eth1/15          Up                    Enabled           100Gx4             1500      -
Eth1/16          Up                    Enabled           100Gx4             1500      -
Eth1/17          Down                  Disabled          Unknown            1500      -
Eth1/18          Down                  Disabled          Unknown            1500      -
Eth1/19          Down                  Disabled          Unknown            1500      -
Eth1/20          Down                  Disabled          Unknown            1500      -
Eth1/21/1        Up                    Enabled           10G                1500      -
Eth1/21/2        Up                    Enabled           10G                1500      -
Eth1/21/3        Up                    Enabled           10G                1500      -
Eth1/21/4        Up                    Enabled           10G                1500      -
Eth1/22          Down                  Disabled          Unknown            1500      -
Eth1/23          Up                    Enabled           10G                1500      -
Eth1/24          Up                    Enabled           10G                1500      -
Eth1/25          Down                  Disabled          Unknown            1500      -
Eth1/26          Down                  Disabled          Unknown            1500      -
Eth1/27          Down                  Disabled          Unknown            1500      -
Eth1/28          Down                  Disabled          Unknown            1500      -
Eth1/29          Down                  Disabled          Unknown            1500      -
Eth1/30          Down                  Disabled          Unknown            1500      -
Eth1/31          Down                  Disabled          Unknown            1500      -
Eth1/32          Down                  Disabled          Unknown            1500      -
```

Related Commands	
Note	• If a high power transceiver (e.g. LR4) is inserted to a port that does not support it, the link does not go up, and the following warning message is displayed: “Warning: High power transceiver is not supported” when running the command “show interfaces ethernet” is run. For more information, please refer to “High Power Transceivers”. • The speed with lane count information refers to the Spectrum-2 and Spectrum-3 systems only.

12.1.7.31 disable interface ethernet traffic-class congestion-control

	disable interface ethernet <inf> traffic-class <tc> congestion-control interface ethernet <inf> disable traffic-class <tc> congestion-control Disables RED/ECN marking for traffic-class queue on ethernet port.	
Syntax Description	inf	Interface number: <slot>/<port>
	tc	Traffic class. Range 0-7
Default	N/A	
Configuration Mode	config	
History	3.8.2000	
Role	admin	

Example	switch (config) # disable interface ethernet 1/1 traffic-class 5 congestion-control switch (config) # interface ethernet 1/1 disable traffic-class 5 congestion-control
Related Commands	show interfaces ethernet 1/1 congestion-control
Notes	The “no interface ethernet <inf> traffic-class <tc> congestion-control” command returns configuration on the port to its default value.

12.1.7.32 disable interface port-channel traffic-class congestion-control

	disable interface port-channel <inf> traffic-class <tc> congestion-control Disables RED/ECN marking for traffic-class queue on LAG port.	
Syntax Description	inf	Interface number. Range: 1-4096
	tc	Traffic class. Range 0-7
Default	N/A	
Configuration Mode	config	
History	3.8.2000	
Role	admin	
Example	switch (config) # disable interface port-channel 15 traffic-class 5 congestion-control switch (config) # interface port-channel 15 disable traffic-class 5 congestion-control	
Related Commands	show interfaces port-channel congestion-control	
Notes	The “no interface port-channel <inf> traffic-class <tc> congestion-control” command returns configuration on the port to its default value.	

12.1.7.33 disable interface mlag-port-channel traffic-class congestion-control

	disable interface mlag-port-channel <inf> traffic-class <tc> congestion-control interface mlag-port-channel <inf> disable traffic-class <tc> congestion-control Disables RED/ECN marking for traffic-class queue on MLAG port.	
Syntax Description	inf	Interface number. Range: 1-1000
	tc	Traffic class. Range 0-7
Default	N/A	
Configuration Mode	config	
History	3.8.2000	
Role	admin	

Example	`switch (config) # disable interface mlag-port-channel 1 traffic-class 5 congestion-control` `switch (config) # interface mlag-port-channel 1 disable traffic-class 5 congestion-control`
Related Commands	show interfaces mlag-port-channel 1/1 congestion-control
Notes	The “no interface mlag-port-channel <inf> traffic-class <tc> congestion-control” command returns configuration on the port to a default value.

12.2 Interface Isolation

Interface isolation provides the ability to group interfaces in sets where traffic from each port is isolated from other interfaces in the group. The isolated interfaces in the group, however, are able to communicate with the interface marked as privileged.

12.2.1 Configuring Isolated Interfaces

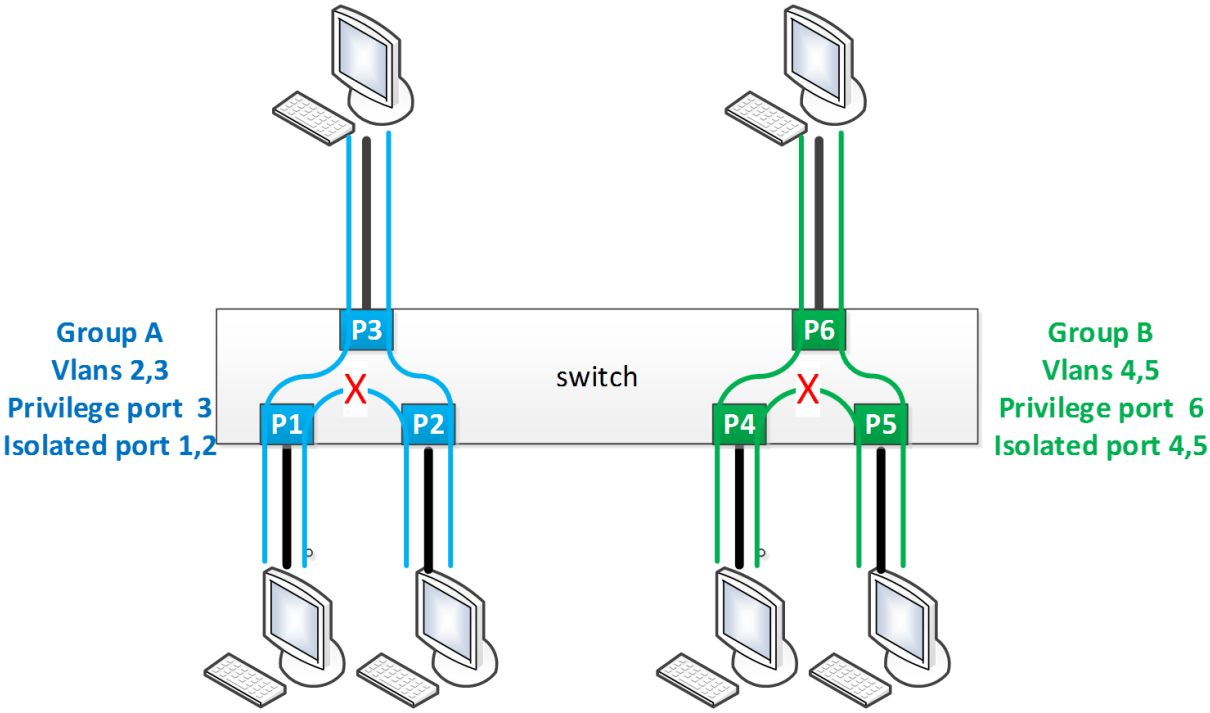

1. Create the VLANs to be used.

```
switch (config) # vlan 2-5
switch (config vlan 2-5) # exit
```

2. Unlock isolation interface protocol.

```
switch (config) # protocol isolation-group
```

3. Create isolation Group A.

```
switch (config) # isolation-group GroupA
```

4. Assign VLANs 2 and 3 to isolation Group A.

```
switch (config isolation-group GroupA) # vlan 2-3
switch (config isolation-group GroupA) # exit
```

5. Create isolation Group B.

```
switch (config) # isolation-group GroupB
```

6. Assign VLANs 4 and 5 to isolation Group B.

```
switch (config isolation-group GroupB) # vlan 4-5
switch (config isolation-group GroupB) # exit
```

7. Set Ethernet interfaces 1-3 to access for VLAN 3.

```
switch (config) # interface ethernet 1/1 switchport access vlan 3
switch (config) # interface ethernet 1/2 switchport access vlan 3
switch (config) # interface ethernet 1/3 switchport access vlan 3
```

8. Isolate Ethernet interfaces 1 and 2 and set Ethernet interfaces 3 as privileged.

```
switch (config) # interface ethernet 1/1-1/2 isolation-group GroupA mode isolated
switch (config) # interface ethernet 1/3 isolation-group GroupA mode privileged
```

9. Enable isolation Group A.

```
(config) # isolation-group GroupA no shutdown
```

10. Set Ethernet interfaces 4-6 to trunk.

```
switch (config) # interface ethernet 1/4 switchport mode trunk
switch (config) # interface ethernet 1/5 switchport mode trunk
switch (config) # interface ethernet 1/6 switchport mode trunk
```

11. Isolate Ethernet interfaces 4 and 5 and set Ethernet interfaces 6 as privileged.

```
switch (config) # interface ethernet 1/4-1/5 isolation-group GroupA mode isolated
switch (config) # interface ethernet 1/6 isolation-group GroupA mode privileged
```

12. Enable isolation Group B.

```
switch (config) # isolation-group GroupB no shutdown
```

13. Verify configuration.

```
switch (config) # show isolation-group
Isolation group: GroupA
State:           Enabled
VLANs:           2, 3
Privileged port: Eth1/3
Isolated ports:  Eth1/1, Eth1/2

Isolation group: GroupB
State:           Enabled
VLANs:           4, 5
Privileged port: Eth1/6
Isolated ports:  Eth1/4, Eth1/5
```

12.2.2 Interface Isolation Commands

12.2.2.1 protocol isolation-group

	protocol isolation-group no protocol isolation-group Enables interface isolation and unlocks further isolation-group commands. The no form of the command disables interface isolation and locks other isolation-group commands.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.6.1002
Example	`switch (config) # protocol isolation-group`
Related Commands	show isolation-group
Notes	• MLAG must be disabled before enabling interface isolation • When disabled, all configuration is lost

12.2.2.2 isolation-group

	isolation-group <name> no isolation-group <name> Creates isolation group. The no form of the command deletes isolation group.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.1002
Example	`switch (config) # isolation-group mygroup` `switch (config isolation-group mygroup) #`
Related Commands	protocol isolation-group show isolation-group
Notes	• The no form of this command deletes the isolation group, removes its attached ports, and the VLANs from the group • Up to 64 isolation groups can be created

12.2.2.3 shutdown

	shutdown no shutdown Disables isolation group. The no form of the command enables isolation group.
Syntax Description	N/A
Default	Disabled

Configuration Mode	config isolation group
History	3.6.1002
Example	`switch (config isolation-group mygroup) # no shutdown`
Related Commands	protocol isolation-group isolation-group show isolation-group
Notes	Enabling isolation groups fails if there are VLANs with ports both inside and outside the group

12.2.2.4 vlan

	vlan <vid> no vlan <vid> Adds a VLAN to isolation group. The no form of the command removes a VLAN from an isolation group.
Syntax Description	N/A
Default	N/A
Configuration Mode	config isolation group
History	3.6.1002
Example	`switch (config isolation-group mygroup) # vlan 10`
Related Commands	protocol isolation-group isolation-group show isolation-group
Notes	• Enabling isolation groups fails if there are VLANs with ports both inside and outside the group • The VLAN must be created before running this command • All interfaces in the VLAN must be attached to only this isolation group • The VLAN added cannot have a respective VLAN interface

12.2.2.5 isolation-group mode

	isolation-group <name> mode {isolated | privileged} no isolation-group <name> mode {isolated | privileged} Adds a VLAN to isolation group. The no form of the command removes a VLAN from an isolation group.	
Syntax Description	name	The isolation group name
	isolated	Configures this interface as isolated
	privileged	Configures this interface as privileged
Default	N/A	
Configuration Mode	config interface ethernet config interface port-channel	
History	3.6.1002	
Example	`switch (config interface ethernet 1/2) # isolation-group mygroup mode privileged`	

Related Commands	protocol isolation-group isolation-group show isolation-group
Notes	

12.2.2.6 show isolation-group

	show isolation-group <name> Displays isolation group information.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	
	3.6.5000	Updated Example
Example	switch (config) # show isolation-group mygroup Isolation group 1: State: Disabled VLANs: N/A Privileged port: N/A Isolated ports: N/A	
Related Commands		
Notes		

12.3 Link Aggregation Group (LAG)

LAG implementation is compliant with 802.1AX standard.

Link Aggregation Group (LAG) protocol describes a network operation in which several same speed links are combined into a single logical entity with the accumulated bandwidth of the originating ports. LAG groups exchange Lag Aggregation Control Protocol (LACP) packets in order to align the functionality between both endpoints of the LAG. To equally send traffic on all LAG links, the switch uses a hash function which can use a set of attributes as key to the hash function.

As many as 32 physical ports can be aggregated on a single LAG.

12.3.1 Configuring Static LAG

1. Create a port-channel entity.

```
switch (config) # interface port-channel 1
switch (config interface port-channel 1) #
```

2. Change back to config mode.

```
switch (config interface port-channel 1) # exit
switch (config) #
```

3. Add a physical port to the LAG.

```
switch (config interface ethernet 1/4) # channel-group 1 mode on
switch (config interface ethernet 1/4) #
```

> If the physical port is operationally up, this port becomes an active member of the aggregation. Consequently, it becomes able to convey traffic.

12.3.2 Configuring Link Aggregation Control Protocol (LACP)

1. Create a port-channel entity.

```
switch (config) # interface port-channel 1
switch (config interface port-channel 1) #
```

2. Change back to config mode.

```
switch (config interface port-channel 1) # exit
switch (config) #
```

3. Enable LACP in the switch.

```
switch (config) # lacp
```

4. Add a physical port to the LAG.

```
switch (config interface ethernet 1/4) # channel-group 1 mode active
```

Or:

```
switch (config interface ethernet 1/4) # channel-group 1 mode passive
```

12.3.3 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community posts:

- HowTo Configure LACP

12.3.4 LAG Commands

12.3.4.1 interface port-channel

	interface port-channel <1-4096>[-<2-4096>] no interface port-channel <1-4096>[-<2-4096>] Creates a LAG and enters the LAG configuration mode. There is an option to create a range of LAG interfaces. The no form of the command deletes the LAG, or range of LAGs.	
Syntax Description	1-4096 / 2-4096	LAG number
Default	N/A	
Configuration Mode	config	
History	3.1.1400	
	3.2.1100	Added range support

Example	`switch (config)# interface port-channel 1`
	`switch (config interface port-channel 1) # exit`
	`switch (config)# interface port-channel 1-10`
	`switch (config interface port-channel 1-10) #`	

Related Commands	show interface port-channel	
Notes	• If a LAG is also an IPL, attempting to delete it without first deleting the IPL is rejected by the management • LAGs have forwarding mode in accordance with the global configuration	

12.3.4.2 lacp

	lacp no lacp Enables LACP in the switch. The no form of the command disables LACP in the switch.
Syntax Description	N/A
Default	LACP is disabled
Configuration Mode	config
History	3.1.1400
Example	`switch (config)# lacp`
Related Commands	
Notes	

12.3.4.3 lacp system-priority

	lacp system-priority <1-65535> no lacp system-priority Configures the LACP system priority. The no form of the command sets the LACP system-priority to default.	
Syntax Description	1-65535	LACP system-priority

Default	32768
Configuration Mode	config
History	3.1.1400
Example	`switch (config)# lacp system-priority 1`
Related Commands	show lacp interfaces port-channel
Notes	Each device that runs LACP has an LACP system priority value. A value between 1 and 65535 can be configured. LACP uses the system priority with the MAC address to form the system ID. When setting the priority, a higher number means a lower priority.

12.3.4.4 lacp (interface)

	lacp {rate fast | port-priority <1-65535>} no lacp {rate fast | port-priority} Configures the LACP interface parameters. The no form of the command sets the LACP interface configuration to default.	
Syntax Description	rate fast	Sets LACP PDUs on the port to be in fast (1 second) or slow rate (30 seconds)
	1-65535	LACP port-priority
Default	rate—slow (30 seconds) port-priority—32768	
Configuration Mode	config interfaces ethernet	
History	3.1.1400	
Example	`switch (config interfaces ethernet 1/7)# lacp rate fast`	
Related Commands		
Notes	Configuring LACP rate (fast or slow) will configure the peer port to send (fast or slow), it does not make any affect on the local port LACP rate.	

12.3.4.5 port-channel load-balance ethernet

	port-channel load-balance ethernet {<method> | [symmetric]} no port-channel load-balance ethernet {<method> | [symmetric]} Configures the port-channel load balancing distribution function method, with symmetric hashing enabled or not. The no form of the command sets the distribution function method to default, or disabling symmetric hashing.		
Syntax Description	method	destination-ip	Destination IP address
		destination-mac	Destination MAC address
		destination-port	Destination UDP/TCP port
		flow-label	IPv6 flow-label field
		l2-protocol	Ethertype field
		l3-protocol	IP protocol field
		ingress-port	Ingress port

		source-destination-ip	Source and destination IP addresses
		source-destination-mac	Source and destination MAC addresses
		source-destination-port	Source and destination UDP/TCP ports
		source-ip	Source IP address
		source-mac	Source MAC address
		source-port	Source UDP/TCP port
		symmetric	Symmetric hashing; bidirectional flows follow same path
	symmetric	Enables symmetric hashing	
Default	source-destination-mac, source-destination-ip, source-destination-port, l3-protocol, l2-protocol, flow-label		
Configuration Mode	config		
History	3.1.1400		
	3.8.1000	Updated syntax	
	3.8.2100	Changed the method options. Modified default LAG HASH to support TCP/UDP ports.	

Example	`switch (config) # port-channel load-balance ethernet ?`
	
	`destination-ip Destination IP address`
	`destination-mac Destination MAC address`
	`destination-port Destination UDP/TCP port`
	`flow-label IPv6 flow-label field`
	`l2-protocol Ethertype field`
	`l3-protocol IP protocol field`
	
	`ingress-port Ingress port`
	`source-destination-ip Source and destination IP addresses`
	`source-destination-mac Source and destination MAC addresses`
	`source-destination-port Source and destination UDP/TCP ports`
	`source-ip Source IP address`
	`source-mac Source MAC address`
	`source-port Source UDP/TCP port`
	`symmetric Symmetric hashing; bidirectional flows follow same path`		

Related Commands	show interface port-channel load-balance		

Notes	• As of 3.8.2100, the default value of port-channel load-balance has been changed from "source-destination-mac" to "source-destination-mac, source-destination-ip, source-destination-port, l3-protocol, l2-protocol, flow-label". This occurs only upon fresh installations or after "reset factory". Upgrading users will retain the old load balancing value and show running-config will indicate this.
	• Several load balance methods can be configured (refer to the example)
	• "ingress-port" and "symmetric" cannot both be set at the same time. The command will be rejected under the following conditions:
	  • 1) "ingress-port" and "symmetric" both appear in the same command.
	  • 2) "ingress-port" is requested while "symmetric" is in force from a previous command. It needs to be cancelled first with "no port-channel load-balance ethernet symmetric".
	  • 3)"symmetric" is requested BY ITSELF while "ingress-port" is in force from a previous command. If "symmetric" is part of a larger list that does not include "ingress-port", the meaning is to exclude "ingress-port" and the command will be accepted.
	• When symmetric is set without other methods: only symmetric hashing can be set while other methods remain unchanged
	• When symmetric is set together with other methods: symmetric hashing is set in parallel with other methods
	• When other methods are set without symmetric: other methods are set, while symmetric hashing remains unchanged

12.3.4.6 channel-group

	channel-group <1-4096> [mode {on | active | passive}] no channel-group Assigns and configures a physical interface to a LAG. The no form of the command removes a physical interface from the port-channel.	
Syntax Description	1-4096	The port channel number
	mode on	Static assignment the port to LAG. LACP will not be enabled on this port.
	mode active/ passive	Dynamic assignment of the port to LAG. LACP will be enabled in either passive or active mode.
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.1.1400	
	3.4.0008	Added a note
	3.6.3640	Added a note
	3.6.4006	Added a note
Example	`switch (config interface ethernet 1/7) # channel-group 1 mode active`	
Related Commands	show interfaces port-channel summary show interfaces port-channel compatibility-parameters show lacp interfaces ethernet	

Notes	• Setting the mode to active/passive is possible only in LACP is enabled • The first port in the LAG decide if the LAG will be static (“on”) or LACP (“active” , “pasive”) • All the ports in the LAG must have the same configuration, determines by the first port added to the LAG. The port with a different configuration will be rejected, for the list of dependencies refer to “show interfaces port-channel compatibility-parameters”. • A physical port may only be part of one channel-group • Added support to check if the forwarding mode of the interface is the same as the forwarding mode of LAG. Error output: *% Channel-group and Ethernet port have different port forwarding mode configuration* • Port cannot be added to port-channel when storm-control is configured on port. Error output: *% Interface * has storm control configuration and can't be added to LAG*

12.3.4.7 lacp-individual enable

	lacp-individual enable [force] no lacp-individual enable [force] Configures the LAG to act with LACP-individual capabilities. The no form of the command disables the LACP-individual capability.	
Syntax Description	force	Toggles the interface after enabling LACP-individual
Default	N/A	
Configuration Mode	config interface port-channel	
History	3.4.1100	
Example	`switch (config interface port-channel 10) # lacp-individual enable force`	
Related Commands		
Notes	If a switch is connected via LAG to a host without LACP capability, running this command on that LAG allows a member port (with the lowest numerical priority value), acting as an individual, to communicate with the host	

12.3.4.8 ip address dhcp

	ip address dhcp no ip address dhcp Enables DHCP on this LAG interface. The no form of the command disables DHCP on this LAG interface.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface port-channel set as router interface
History	3.4.2008
Example	`switch (config interface port channel 10) # ip address dhcp`
Related Commands	interface port-channel show interface port-channel
Notes	

12.3.4.9 show lacp counters

	show lacp counters Displays the LACP PDUs counters.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config interface port-channel set as router interface	
History	3.1.1400	
	3.6.6000	Updated example

Example

```
switch (config) # show lacp counters
VRF Name: default

Port-channel 5:
  ---------------------------------------------------------------------------------
  LACPDUs  Marker  Marker  Marker Rsp  Marker Rsp  LACPDUs  LACPDUs  Illegal  Unknown
  Port     Sent    Recv    Sent        Recv        Sent     Recv
  ---------------------------------------------------------------------------------
  1/12     0       0       0           0           0        0        0        0
  1/11     0       0       0           0           0        0        0        0
  1/10     0       0       0           0           0        0        0        0
```

Related Commands	interface port-channel show interface port-channel
Notes	

12.3.4.10 show lacp interfaces ethernet

	show lacp interface ethernet <inf> Displays the LACP interface configuration and status.	
Syntax Description	inf	Interface number (e.g., “1/1”)
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.1400	
	3.6.6102	Updated example

Example

```
switch (config) # show lacp interfaces ethernet 1/1
Port: 1/1
Port State: Down
Channel Group: 1
Pseudo port-channel: Po1
LACP port-priority: 32768
LACP Rate: Slow
LACP Activity: Active
LACP Timeout: Short
Aggregation State: Aggregation, Defaulted,
------------------------------------------------------------
                  LACP Port  Admin   Oper  Port       Port
Port     State    Priority   Key     Key   Number     State
------------------------------------------------------------
1/1      Down     32768      13826   13826 0x1        0x0
```

Related Commands	
Notes	

12.3.4.11 show lacp interfaces neighbor

	show lacp interfaces neighbor Displays the LACP interface neighbor status.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.1400	
	3.4.0000	Updated example

Example

```
switch (config) # show lacp interfaces neighbor
Flags:
A - Device is in Active mode
P - Device is in Passive mode

Channel group 1 neighbors

Port 1/4
----------
Partner System ID               : 00:00:00:00:00:00
Flags                           : A
LACP Partner Port Priority      : 0
LACP Partner Oper Key           : 0
LACP Partner Port State         : 0x0

Port State Flags Decode
-----------------------
Activity : Active
Aggregation State : Aggregation, Sync, Collecting, Distributing

MLAG channel group 25 neighbors

Port 1/49
----------
Partner System ID               : 00:02:c9:fa:c4:c0
Flags                           : A
LACP Partner Port Priority      : 255
LACP Partner Oper Key           : 33
LACP Partner Port State         : 0xbc

Port State Flags Decode
-----------------------
Activity : Active
Aggregation State : Aggregation, Sync, Collecting, Distributing,

MLAG channel group 28 neighbors

Port 1/51
----------
Partner System ID               : f4:52:14:10:d8:f1
Flags                           : A
LACP Partner Port Priority      : 255
LACP Partner Oper Key           : 33
LACP Partner Port State         : 0xbc

Port State Flags Decode
-----------------------
Activity : Active
Aggregation State : Aggregation, Sync, Collecting, Distributing,
```

Related Commands	
Notes	

12.3.4.12 show lacp

	show lacp Displays the LACP global parameters.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.4.0000
Example	`switch (config) # show lacp` `Port-channel Module Admin Status is enabled`
Related Commands	
Notes	

12.3.4.13 show lacp interfaces system-identifier

	show lacp interfaces {mlag-port-channel | port-channel} <instance> system-identifier Displays the system identifier of LACP.	
Syntax Description	instance	LAG or MLAG instance
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.0000	

Example	`switch (config)# show lacp interfaces port-channel 2 system-identifier`
	`Priority: 12345`
	`MAC: 00:02:c9:ac:2a:60`	

Related Commands		
Notes		

12.3.4.14 show interfaces port-channel

	show interfaces port-channel <port-channel> Displays LAG configuration properties.	
Syntax Description	port-channel	LAG interface whose properties to display
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4000	
	3.4.1100	Updated example
	3.6.1002	Added “error packets” counter to Tx
	3.6.5000	Updated example with telemetry
	3.6.8008	Updated example
	3.7.1000	Updated example

	3.9.1000	Updated example

Example

```
switch (config) # show interfaces port-channel 10

Po10:
  Admin state         : Enabled
  Operational state   : Down
  Description         : N/A
  Mac address         : N/A
  MTU                 : 1500 bytes (Maximum packet size 1522 bytes)
  lacp-individual mode: Disabled
  Flow-control        : receive off send off
  Actual speed        : N/A
  Width reduction mode: Not supported
  Switchport mode     : access
  MAC learning mode   : Enabled
  Forwarding mode     : inherited cut-through
  FCS Ingress         : Enabled CRC check
  FCS Egress          : Disabled CRC recalculate
  FCS Timestamping    : Enabled

  Telemetry sampling: Disabled   TCs: N/A
    Telemetry threshold: Disabled        TCs: N/A
    Telemetry threshold level: N/A

  Last clearing of "show interface" counters: Never
  60 seconds ingress rate                   : 0 bits/sec, 0 bytes/sec, 0 packets/sec
  60 seconds egress rate                    : 0 bits/sec, 0 bytes/sec, 0 packets/sec

  Rx:
    0                    packets
    0                    unicast packets
    0                    multicast packets
    0                    broadcast packets
    0                    bytes
    0                    discard packets
    0                    error packets
    0                    fcs errors
    0                    undersize packets
    0                    oversize packets
    0                    pause packets
    0                    unknown control opcode
    0                    symbol errors
    0                    discard packets by storm control

  Tx:
    0                    packets
    0                    unicast packets
    0                    multicast packets
    0                    broadcast packets
    0                    bytes
    0                    discard packets
    0                    error packets
    0                    hoq discard packets
```

Related Commands	
Notes	

12.3.4.15 show interfaces port-channel counters

	show interfaces port-channel <port-channel> counters Displays the extended counters for the interface.	
Syntax Description	port-channel	LAG interface whose properties to display.
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	

	3.9.1000	Added ability to use a range of ports
Example	switch (config) # show interfaces port-channel 2-3 counters Po2: Rx: 0 packets 0 unicast packets 0 multicast packets 0 broadcast packets 0 bytes 0 packets of 64 bytes 0 packets of 65-127 bytes 0 packets of 128-255 bytes 0 packets of 256-511 bytes 0 packets of 512-1023 bytes 0 packets of 1024-1518 bytes 0 packets Jumbo 0 error packets 0 discard packets 0 fcs errors 0 undersize packets 0 oversize packets 0 pause packets 0 unknown control opcode 0 symbol errors Tx 1000000 packets 0 unicast packets 1000000 multicast packets 0 broadcast packets 1505000000 bytes 1000000 error packets 0 discard packets 0 pause packets 0 ECN marked packets Po3: . . .	
Related Commands		
Notes	As of version 3.9.1000, the "port-channel" attribute is optional. If nothing is selected, information for all ports will be displayed	

12.3.4.16 show interfaces port-channel compatibility-parameters

	show interfaces port-channel compatibility-parameters Displays LAG parameters.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4000	
	3.6.3640	Added “forwarding mode” as compatibility parameter to output
	3.6.6000	Updated example
	3.6.8008	Updated example

Example	switch (config) # show interfaces port-channel compatibility-parameters
	
	Compatibility-parameters:
	  * Port-mode
	  * Speed
	  * MTU
	  * Forwarding mode
	  * Flow Control
	  * Access VLAN
	  * Allowed VLAN list
	  * Flowcontrol & PFC
	  * Channel-group mode
	  * QoS parameters
	  * MAC learning disable
	
	Static configuration on the port should be removed:
	  * ACL port binding
	  * Static mrouter
	  * sflow
	  * OpenFlow
	  * port mirroring local analyzer port
	  * Static mac address

Related Commands	
Notes	

12.3.4.17 show interfaces port-channel load-balance

	show interfaces port-channel load-balance Displays the type of load-balancing in use for LAGs.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.4000
Example	switch (config) # show interfaces port-channel load-balance source-destination-mac
Related Commands	port-channel load-balance ethernet ?
Notes	

12.3.4.18 show interfaces port-channel summary

	show interfaces port-channel summary Displays a summary for LAG interfaces.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.1400	
	3.4.1100	Updated example
Example		

```
switch (config) # show interfaces port-channel summary

Flags: D - Down, U - Up, P - Up in port-channel (members)
       S - Suspend in port-channel (members), I - Individual

--------------------------------------------------------------------------------
Group Port-      Type       Member Ports
Channel
--------------------------------------------------------------------------------
1 Po2(U)         LACP       Eth1/58(D) Eth1/59(I) Eth1/60(S)
2 Po5(D)         LACP       Eth1/1(S) Eth1/33(I)
3 Po10(U)        LACP       Eth1/49(P) Eth1/50(P) Eth1/51(S) Eth1/52(S)
```

Related Commands	
Notes	

12.4 Link Layer Discovery Protocol (LLDP)

The Link Layer Discovery Protocol (LLDP) is a vendor-neutral Link Layer protocol in the Internet Protocol Suite used by network devices for advertising their identity, capabilities, and neighbors on a IEEE 802 LAN. The protocol is formally defined in IEEE 802.1AB. From version 3.8.2000, LLDP is now enabled by default.

12.4.1 Configuring LLDP

1. Enable LLDP globally on the switch.

```
switch (config) # lldp
```

2. Enable LLDP per interface.

```
switch (config interface ethernet 1/1) # lldp receive
switch (config interface ethernet 1/1) # lldp transmit
```

3. Display LLDP local information.

```
switch (config) # show lldp local

LLDP is Enabled

Local global configuration
Chassis sub type: macAddress (4)
Chassis id: 00:11:22:33:44:55
System Name: "switch-111111"
System Description: my-system-description
Supported capabilities: B
Supported capabilities enabled: B
```

4. Display LLDP remote information.

```
switch (config)# show lldp interfaces ethernet 1/1 remote

Ethernet 1/1
Remote Index: 1
Remote chassis id: 00:11:22:33:44:55 ; chassis id subtype: mac
Remote port-id: ethenret 1/2; port id subtype: local
Remote port description: ethernet 1/2
Remote system name: remote-system
Remote system description: remote-system-description
Remote system capabilities supported: B ; B
```

12.4.2 DCBX

Data Center Bridging (DCB) is an enabler for running the Ethernet network with lossless connectivity using priority-based flow control and enhanced transmission selection. DCBX (exchange) complements the DCB implementation by offering a dynamic protocol that communicates DCB attributes between peering endpoint. NVIDIA Onyx supports two versions of DCBX TLVs running on top of LLDP:

- DCBX IEEE
- DCBX CEE

By default DCBX IEEE is enabled when LLDP is enabled. LLDP is enabled by default.

12.4.3 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community posts:

- DCBX Versions and Support on Ethernet Switches
- LLDP DCBX Packet Format Examples IEEE and CEE (Wireshark)

12.4.4 LLDP Commands

12.4.4.1 lldp

	lldp no lldp Enables LLDP globally. The no form of the command disables the LLDP.	
Syntax Description	N/A	
Default	Enabled	
Configuration Mode	config	
History	3.2.0300	
	3.8.2000	Changed default from "disabled" to "enabled"
Example	`switch (config)# lldp`	
Related Commands	show lldp local	
Notes		

12.4.4.2 lldp reinit

	lldp reinit <seconds> no lldp reinit Sets the delay in seconds from enabling the LLDP on the port until re-initialization will be attempted. The no form of the command sets the parameter to default.	
Syntax Description	seconds	1-10

Default	2
Configuration Mode	config
History	3.2.0300
Example	`switch (config)# lldp reinit 10`
Related Commands	show lldp timers
Notes	

12.4.4.3 lldp timer

	lldp timer <seconds> no lldp timer Sets the LLDP interval at which LLDP frames are transmitted. (lldpMessageTxInterval). The no form of the command sets the parameter to default.	
Syntax Description	seconds	5-32768
Default	30	
Configuration Mode	config	
History	3.2.0300	
Example	`switch (config)# lldp timer 10`	
Related Commands	show lldp timers	
Notes		

12.4.4.4 lldp tx-delay

	lldp tx-delay <seconds> no lldp tx-delay Indicates the delay in seconds between successive LLDP frame transmissions. The no form of the command sets the parameter to default.	
Syntax Description	seconds	1-8192
Default	2	
Configuration Mode	config	
History	3.2.0300	
Example	`switch (config)# lldp tx-delay 10`	
Related Commands	show lldp timers	
Notes	The recommended value for the tx-delay is set by the following formula: 1 <= lldp tx-delay <= (0.25 * lldp timer)	

12.4.4.5 lldp tx-hold-multiplier

	lldp tx-hold-multiplier <seconds> no lldp tx-hold-multiplier The time-to-live value expressed as a multiple of the lldpMessageTxInterval object. The no form of the command sets the parameter to default.	
Syntax Description	seconds	1-8192
Default	2	
Configuration Mode	config	
History	3.2.0300	
Example	`switch (config)# lldp tx-hold-multiplier 10`	
Related Commands	show lldp timers	
Notes	The actual time-to-live value used in LLDP frames, can be expressed by the following formula: TTL = min(65535, (lldpMessageTxInterval * lldpMessageTxHoldMultiplier)). For example, if the value of lldpMessageTxInterval is 30, and the value of lldpMessageTxHoldMultiplier is 4, then the value 120 is encoded in the TTL field in the LLDP header.	

12.4.4.6 lldp (interface)

	lldp {receive | transmit} no lldp {receive | transmit} Enables LLDP receive or transmit capabilities. The no form of the command disables LLDP receive or transmit capabilities.	
Syntax Description	med-tlv-select	Enables LLDP media TLVs.
	receive	Enables LLDP receive on this port.
	tlv-select	Enables LLDP TLVs.
	transmit	Enables LLDP transmit on this port.
Default	Enabled for receive and transmit	
Configuration Mode	config interface ethernet	
History	3.2.0300	
Example	`switch (config interface ethernet 1/1)# lldp receive`	
Related Commands	show lldp interface	
Notes	The LLDP is disabled by default (globally)	

12.4.4.7 lldp tlv-select

	lldp tlv-select {[dcbx] [dcbx-cee] [port-description] [sys-name] [sys-description] [sys-capabilities] [management-address] [none] all} Sets the LLDP basic TLVs to be transmitted on this port.	
Syntax Description	dcbx	Enables LLDP-DCBX TLVs
	dcbx-cee	Enables LLDP-DCBX CEE TLVs
	port-description	LLDP port description TLV

	sys-name	LLDP system name TLV
	sys-description	LLDP system description TLV
	sys-capabilities	LLDP system capabilities TLV
	management-address	LLDP management address TLV
	all	all above TLVs
	none	None of the above TLVs
Default	all	
Configuration Mode	config interface ethernet	
History	3.2.0300	
	3.3.0000	Added “none” parameter
	3.3.4302	Added “dcbx” parameter
	3.3.4402	Added “dcbx-cee” parameter
Example	`switch (config interface ethernet 1/1)# lldp tlv-select port-description sys-name`	
Related Commands	show lldp interface	
Notes	The management address is chosen according to the following criteria where 1 takes priority over 2, and 2 takes priority over 3: 1. Smallest IP address of mgmt0 2. Smallest IP address of mgmt1 3. First primary address of all non-management interfaces	

12.4.4.8 lldp med-tlv-select

	lldp med-tlv-select {all | media-capability | network-policy | none} Configures LLDP media TLV attributes.	
Syntax Description	all	Enables all LLDP media TLVs
	media-capabilities	Enables Media Capabilities TLV
	network-policy	Enables Network-Policy TLV
	none	Disables all LLDP media TLVs
Default	Disabled	
Configuration Mode	config interface ethernet	
History	3.6.1002	
Example	`switch (config interface ethernet 1/1)# lldp med-tlv-select all`	
Related Commands	show lldp interface	
Notes		

12.4.4.9 dcb application-priority

	dcb application-priority <selector> <protocol> <priority> Adds an application to the application priority table.

Syntax Description	selector	Protocol type: ethertype
	protocol	Protocol field in hexadecimal notation (e.g. '0x8906' for FCoE, '0x8914' for FIP)
	priority	Range: 0-7
Default	No applications are available. The table is empty.	
Configuration Mode	config	
History	3.3.4200	
Example	`switch (config-if)# dcb application-priority ethertype 0x8906`	
Related Commands	show lldp interface	
Notes		

12.4.4.10 clear lldp counters

	clear lldp counters [<Device | Port>] Clears LLDP counters for all ports or for a specific port.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.4006
Example	`switch (config) # clear lldp counters` `switch (config) # clear lldp counters 1/1`
Related Commands	
Notes	

12.4.4.11 show lldp local

	show lldp local Displays LLDP local information.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.2.0300

Example	`switch (config)# show lldp local`
	`LLDP is Enabled`
	`Local global configuration`
	`Chassis sub type: macAddress (4)`
	`Chassis id: 0002C9030046AF00`
	`System Name: my-switch`
	`System Description: SN2100`
	`Supported capabilities: B,R`
	`Supported capabilities enabled: B`

Related Commands	
Notes	

12.4.4.12 show lldp interfaces

	show lldp interfaces [ethernet <inf> [med-cap | remote]] Displays LLDP remote interface table information.	
Syntax Description	inf	Local interface number (e.g. 1/1)
	med-cap	Displays local port media capabilities information
	remote	Displays LLDP Ethernet remote configuration & status
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.0300	
	3.3.4200	Updated example
	3.6.1002	Updated example
Example		

```
switch (config)# show lldp interfaces
TLV flags:
PD: port-description, SN: sys-name, SD: sys-description, SC: sys-capabilities, MA: management-address
ETS-C: ETS-Configuration, ETS-R: ETS-Recommendation, AP: Application Priority, PFC: Priority Flow Control
CEE: Converged Enhanced Ethernet DCBX version
MED-CAP: Media Capabilities
MED-NWP: MED-Network Policy

Interface Receive   Transmit  TLVs
-----------------------------------------------------------------------------------------
Eth1/1    Enabled   Enabled   PD, SD
Eth1/2    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R
Eth1/3    Disabled  Disabled  PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R, MED-NWP
Eth1/4    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R, MED-CAP, MED-NWP
Eth1/5    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R
Eth1/6    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R
Eth1/7    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R
```

Related Commands	
Notes	

12.4.4.13 show lldp remote

	show lldp remote Displays LLDP remote information (remote device id, remote port id, remote system name).
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.3004
Example	

```
switch (config)# show lldp remote
-------------------------------------------------------------------------------
Local Interface      Device ID            Port ID              System Name
-------------------------------------------------------------------------------
Eth1/4               e4:1d:2d:a5:f3:35    e4:1d:2d:a5:f3:35    Not Advertised
Eth1/10              e4:1d:2d:44:65:00    Eth1/10              switch108
Eth1/11              e4:1d:2d:44:65:00    Eth1/11              switch108
Eth1/12              e4:1d:2d:44:65:00    Eth1/12              switch108
Eth1/13              e4:1d:2d:44:65:00    Eth1/13              switch108
Eth1/14              e4:1d:2d:44:65:00    Eth1/14              switch108
Eth1/15              e4:1d:2d:44:65:00    Eth1/15              switch108
Eth1/16              e4:1d:2d:44:65:00    Eth1/16              switch108
Eth1/17              e4:1d:2d:44:65:00    Eth1/17              switch108
Eth1/18              e4:1d:2d:44:65:00    Eth1/18              switch108
Eth1/19              e4:1d:2d:44:65:00    Eth1/19              switch108
Eth1/20              e4:1d:2d:44:65:00    Eth1/20              switch108
Eth1/21              e4:1d:2d:44:65:00    Eth1/21              switch108
Eth1/22              e4:1d:2d:44:65:00    Eth1/22              switch108
Eth1/23              e4:1d:2d:44:65:00    Eth1/23              switch108
Eth1/24              e4:1d:2d:44:65:00    Eth1/24              switch108
Eth1/25              e4:1d:2d:44:65:00    Eth1/25              switch108
Eth1/26              e4:1d:2d:44:65:00    Eth1/26              switch108
Eth1/31              e4:1d:2d:44:65:00    Eth1/31              switch108
Eth1/32              e4:1d:2d:44:65:00    Eth1/32              switch108
```

Related Commands	
Notes	

12.4.4.14 show lldp statistics

	show lldp statistics [<inf>] Displays LLDP interface statistics.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.2.0300
Example	

```
switch (config)# show lldp statistics
-------------------------------------------------------------------------------------------------------
Interface    Frames       In           In           TLVs         TLVs           Ageout       Out
             Discarded    Errors       Total        Discarded    Unrecognized                Frames
-------------------------------------------------------------------------------------------------------
Eth1/1       0            0            0            0            0              0            0
Eth1/2       0            0            20           0            40             0            5
Eth1/3       16           0            16           0            0              0            0
Eth1/4       0            0            15           0            30             0            5
Eth1/5       0            0            15           0            30             0            5
Eth1/6       0            0            0            0            0              0            0
Eth1/7       0            0            0            0            0              0            0
Eth1/8       0            0            0            0            0              0            0
Eth1/9       0            0            0            0            0              0            0
Eth1/10      0            0            5            0            15             0            5
Eth1/12      0            0            5            0            15             0            5
Eth1/13      0            0            5            0            15             0            5
Eth1/14      0            0            0            0            0              0            0
Eth1/15      0            0            6            0            18             0            5
Eth1/16      0            0            5            0            15             0            6
```

Related Commands	
Notes	

12.4.4.15 show lldp statistics global

	show lldp statistics global Displays LLDP global statistics.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.2.0300

Example	`switch (config)# show lldp timers`
	`Remote Table Last Change Time : 10300`
	`Remote Table Inserts : 5`
	`Remote Table Deletes : 0`
	`Remote Table Drops : 0`
	`Remote Table Ageouts : 0`

Related Commands	
Notes	

12.4.4.16 show lldp timers

	show lldp timers Displays LLDP timers configuration
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.2.0300

Example	`switch (config)# show lldp timers`
	`msg-tx-interval :30`
	`tx-delay :2`
	`tx-hold :4`
	`tx-reinit-delay :2`

Related Commands	
Notes	

12.4.4.17 show dcb application-priority

	show dcb application-priority Displays application priority admin table.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.4200

Example	switch (config)# show dcb application-priority ----------------------------------- Selector Protocol Priority ----------------------------------- Ethertype 0x8906 3 Ethertype 0x8914 3
Related Commands	
Notes	

12.5 VLANs

A Virtual Local Area Network (VLAN) is an L2 segment of the network which defines a broadcast domain and is identified by a tag added to all Ethernet frames running within the domain. This tag is called a VLAN ID (VID) and can be assigned a value of 1-4094.

Each port can have a switch mode of either:

- Access – access port is a port connected to a host. It can accept only untagged frames, and assigns them a default configured VLAN (Port VLAN ID). On egress, traffic sent from the access port is untagged.
- Access-dcb – receives ingress untagged traffic but sends egress priority tag (VLAN ID = 0)
- Hybrid – hybrid port is a port connected to either switches or hosts. It can receive both tagged and untagged frames and assigns untagged frames a default configured VLAN (Port VLAN ID). It receives tagged frames with VLANs of which the port is a member (these VLANs' names are allowed). On egress, traffic of allowed VLANs sent from the Hybrid port is sent tagged, while traffic sent with PVID is untagged.
- Trunk – trunk port is a port connecting 2 switches. It accepts only tagged frames with VLANs of which the port is a member. On egress, traffic sent from the Trunk port is tagged. By default, a Trunk port is, automatically, a member on all current VLANs.

12.5.1 Configuring Access Mode and Assigning Port VLAN ID (PVID)

1. Create a VLAN.

```
switch (config) # vlan 6
switch (config vlan 6) #
```

2. Change back to config mode.

```
switch (config vlan 6) # exit
switch (config) #
```

3. Enter the interface configuration mode.

```
switch (config) # interface ethernet 1/22
switch (config interface ethernet 1/22) #
```

4. From within the interface context, configure the interface mode to Access.

```
switch (config interface ethernet 1/22) # switchport mode access
```

5. From within the interface context, configure the Access VLAN membership.

```
switch (config interface ethernet 1/22) # switchport access vlan 6
```

12.5.2 Configuring Hybrid Mode and Assigning Port VLAN ID (PVID)

1. Create a VLAN.

```
switch (config) # vlan 6
switch (config vlan 6) #
```

2. Change back to config mode.

```
switch (config vlan 6) # exit
switch (config) #
```

3. Enter the interface configuration mode.

```
switch (config) # interface ethernet 1/22
switch (config interface ethernet 1/22) #
```

4. From within the interface context, configure the interface mode to Access.

```
switch (config interface ethernet 1/22) # switchport mode hybrid
switch (config interface ethernet 1/22) #
```

5. From within the interface context, configure the Access VLAN membership.

```
switch (config interface ethernet 1/22) # switchport access vlan 6
```

12.5.3 Configuring Trunk Mode VLAN Membership

1. Create a VLAN.

```
switch (config) # vlan 10
switch (config vlan 10) #
```

2. Change back to config mode.

```
switch (config vlan 10) # exit
switch (config) #
```

3. Enter the interface configuration mode.

```
switch (config) # interface ethernet 1/35
switch (config interface ethernet 1/35) #
```

4. From within the interface context, configure the interface mode to Trunk.

```
switch (config interface ethernet 1/35) # switchport mode trunk
```

12.5.4 Configuring Hybrid Mode VLAN Membership

1. Create a VLAN.

```
switch (config) # vlan 10
switch (config vlan 10) #
```

2. Change back to config mode.

```
switch (config vlan 10) # exit
switch (config) #
```

3. Enter the interface configuration mode.

```
switch (config) # interface ethernet 1/35
switch (config interface ethernet 1/35) #
```

4. From within the interface context, configure the interface mode to Hybrid.

```
switch (config interface ethernet 1/35) # switchport mode hybrid
switch (config interface ethernet 1/35) #
```

5. From within the interface context, configure the allowed VLAN membership.

```
switch (config interface ethernet 1/35) # switchport hybrid allowed-vlan add 10
switch (config interface ethernet 1/35) #
```

12.5.5 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community post:

- Howto Configure Switch Port Types with NVIDIA Onyx

12.5.6 VLAN Commands

12.5.6.1 vlan

	vlan {<vlan-id> | <vlan-range>} no vlan {<vlan-id> | <vlan-range>} Creates a VLAN or range of VLANs, and enters a VLAN context. The no form of the command deletes the VLAN or VLAN range.	
Syntax Description	vlan-id	Range: 1-4094
	vlan-range	Any range of VLANs
Default	VLAN 1 is enabled by default	
Configuration Mode	config	
History	3.1.1400	
Example	switch (config) # vlan 10 switch (config vlan 10) #	

Related Commands	show vlan switchport mode switchport [trunk | hybrid] allowed-vlan
Notes	Interfaces are not added automatically to VLAN unless configured with trunk or hybrid mode with “all” option turned on.

12.5.6.2 name

	name <vlan-name> no name Adds VLAN name. The no form of the command deletes the VLAN name.	
Syntax Description	vlan-name	40-character long string
Default	No name available	
Configuration Mode	config vlan	
History	3.1.1400	
Example	`switch (config vlan 10) # name my-vlan-name`	
Related Commands	show vlan switchport mode switchport [trunk | hybrid] allowed-vlan	
Notes	Name can not be configured for a range of VLANs.	

12.5.6.3 show vlan

	show vlan [id <vlan-id>] Displays the VLAN table.	
Syntax Description	vlan-id	1-4094
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.1400	
Example	*switch (config vlan*	
Related Commands	show vlan switchport mode switchport [trunk | hybrid] allowed-vlan vlan	
Notes		

12.5.6.4 switchport mode

	switchport mode {access | dot1q-tunnel | trunk | hybrid | access-dcb} no switchport mode Sets the switch port mode. The no form of the command sets the switch port mode to access.

Syntax Description	access	Untagged port. 802.1q tagged traffic are filtered. Egress traffic is untagged.
	dot1q-tunnel	Allows both tagged and untagged ingress Ethernet packets. Egress packets are tagged with a second VLAN (802.1Q) header.
	trunk	802.1q tagged port, untagged traffic is filtered.
	hybrid	Both 802.1q tagged and untagged traffic is allowed on the port.
	access-dcb	Untagged port, egress traffic is priority tagged.
Default	access	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.1400	
	3.3.4500	Added MPO configuration mode
	3.4.3000	Added dot1q-tunnel parameter
	3.6.6000	Added ability to switchport mode for a range of interfaces
Example	`switch (config) # interface ethernet 1/7` `switch (config interface ethernet 1/7) # switchport mode access`	
Related Commands	show vlan show interfaces switchport switchport access vlan switchport [trunk | hybrid] allowed-vlan switchport dot1q-tunnel qos-mode vlan	
Notes	Switchport mode may be configured for a range of interfaces (interface <inf-type> <id-range> switchport mode <type>)	

12.5.6.5 switchport dot1q-tunnel qos-mode

	switchport dot1q-tunnel qos-mode {pipe | uniform} no switchport dot1q-tunnel qos-mode Assigns QoS to the service provider's traffic. The no form of the command resets the parameter value to its default.	
Syntax Description	pipe	Gives the service provider's traffic QoS 0
	uniform	Gives the service provider's traffic the same QoS as the customer's traffic
Default	pipe	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.4.3000	
Role	admin	
Example	`switch (config interface ethernet 1/1) # switchport dot1q-tunnel qos-mode uniform`	

Related Commands	show vlan show interfaces switchport switchport access vlan switchport [trunk | hybrid] allowed-vlan vlan
Notes	

12.5.6.6 switchport access

	switchport access vlan <vlan-id> no switchport access vlan switchport access none (hybrid mode only) Configures the port access VLAN. The no form of the command sets the port access VLAN to 1. The none clause of the command removes access VLAN membership from the port, thus disallowing untagged traffic on this port. This is commonly used for fast transition from hybrid switchport to trunk-like switchport and vice versa.	
Syntax Description	vlan-id	1-4094
Default	1	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.1400	
	3.2.0500	Format change (removed hybrid and access-dcb options). Previous command format was: “switchport {hybrid | access-dcb | access} vlan <vlan-id>”.
	3.3.4500	Added MPO configuration mode.
	3.6.6000	Added ability to configure VLAN ID for a range of interfaces.
	3.7.1100	Updated command syntax & notes.
Example	`switch (config interface ethernet 1/7) # switchport access vlan 10`	
Related Commands	show vlan show interfaces switchport switchport mode switchport [trunk | hybrid] allowed-vlan vlan	
Note	• This command is not applicable for interfaces with port mode trunk • Only one option (“access”, “access-dcb” or “hybrid”) is possible to configure on the port, depending on the switchport mode of the port • Access VLAN ID may be configured to a range of interfaces (interface <inf-type> <id-range> switchport access vlan <vlan-ID>) • This command is not applicable for interfaces with port mode trunk • In hybrid mode, access vlan is optional. Alternatively, use “access none” in order to disable access vlan. In this case, all incoming untagged traffic will be dropped.	

12.5.6.7 switchport {hybrid, trunk} allowed-vlan

	switchport {hybrid, trunk} allowed-vlan {<vlan> | add <vlan> | remove <vlan> all | except <vlan> | none} Sets the port allowed VLANs.

Syntax Description	vlan	VLAN ID (1-4094) or VLAN range
	add	Adds VLAN or range of VLANs
	remove	Removes VLANs or range of VLANs
	all	Adds all VLANs in available in the VLAN table" New VLANs added to the VLAN table are added automatically
	except	Adds all VLANs expect this VLAN or VLAN range
	none	Removes all VLANs
Default	N/A	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.1400	
Example	`switch (config interface ethernet 1/7) # switchport hybrid allowed-vlan all`	
Related Commands	show vlan show interfaces switchport switchport access vlan switchport mode vlan	
Note	• This command is not applicable for interfaces with port mode access or access-dcb • In order for the parameter “hybrid” or “trunk” to be available, the switchport mode on the interface must be configured to either hybrid or trunk respectively	

12.5.6.8 switchport voice

	switchport voice vlan <vlan-id> no switchport voice vlan Configures voice VLAN for the interface. The no form of the command disables voice VLAN.	
Syntax Description	vlan-id	1-4094
Default	Disabled	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.6.1002	
Example	`switch (config interface ethernet 1/7) # switchport voice vlan 10`	
Related Commands	lldp med-tlv-select show vlan show interfaces switchport switchport mode switchport [trunk | hybrid] allowed-vlan vlan	
Note		

12.5.6.9 show interfaces switchport

	show interfaces [<if>] switchport Displays all interface switch port configurations.	
Syntax Description	if	Possible interface types: • ethernet <slot/port> • port-channel <lag-id> • mlag-port-channel <id>
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.1400	
	3.6.6102	Added ability to filter by specific interfaces and updated Example
Example		

```
switch (config) # show interfaces switchport

-------------------------------------------------------------------
Interface         Mode           Access vlan          Allowed vlans
-------------------------------------------------------------------
Eth1/1            access         1
Eth1/2            access         1
Eth1/6            access         1
```

Related Commands	show vlan switchport access vlan switchport mode vlan
Notes	This command can accept an explicit interface or interface range (displays information only for available interfaces)

12.6 Voice VLAN

Voice VLAN allows configuring a port to provide QoS to voice and data traffic in a scenario where a terminal is connected to an IP phone which is in turn connected to the port on the switch. The IP phone bridges the data traffic from the terminal into the switch port. Any voice traffic from the IP phone is also sent to the same port with no differentiation. Therefore it is in the administrator's interest to provide different QoS to the voice traffic and the data traffic by placing the voice traffic on a different VLAN from the data traffic.

This can be achieved by configuring a voice VLAN on the desired switch port using LLDP-MED TLVs. Media Endpoint Discovery (MED) TLVs allow the switch to apply certain policies by informing the remote media device to configure itself using different TLV.

In this use-case scenario we employ the use of the network policy TLV, which is defined as per TIA-TR41. The network policy TLV can be used to inform a specific VLAN to use for an application stream.

The OS allows the user to configure the VLAN for voice traffic. In the following figure, the user configures a voice VLAN of 25 and the switch port has a PVID of 50. Therefore all the voice traffic is switched onto VLAN 25 and the untagged packets from the terminal are switched into VLAN 50.

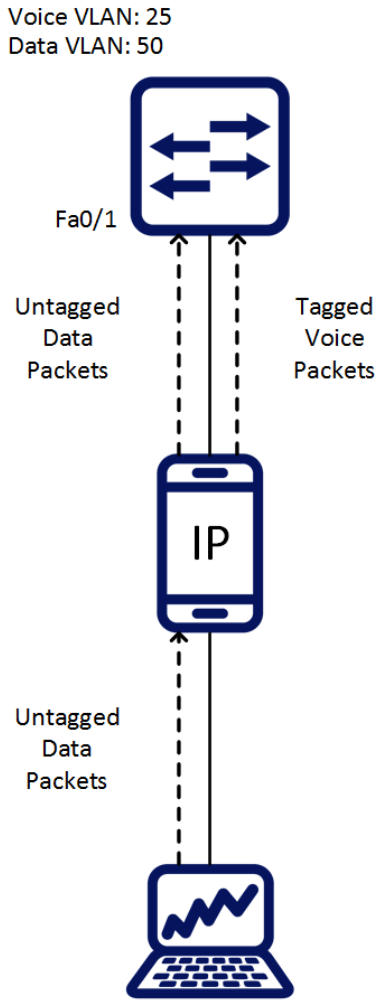

12.6.1 Configuring Voice VLAN

To configure LLDP-MED TLV, run the following:

```
switch (config) # interface ethernet 1/4
switch (config interface ethernet 1/4) # lldp med-tlv-select media-capabilities
switch (config interface ethernet 1/4) # lldp med-tlv-select network-policy
switch (config interface ethernet 1/4) # lldp med-tlv-select all
```

To verify LLDP-MED TLV configuration, run the following:

```
switch (config) # show lldp interface
TLV flags:
PD: port-description, SN: sys-name, SD: sys-description, SC: sys-capabilities, MA: management-address
ETS-C: ETS-Configuration, ETS-R: ETS-Recommendation, AP: Application Priority, PFC: Priority Flow Control
CEE: Converged Enhanced Ethernet DCBX version
MED-CAP: Media Capabilities
MED-NWP: MED-Network Policy

Interface Receive   Transmit  TLVs
-------------------------------------------------------------------------------------
Eth1/1    Enabled   Enabled   PD, SD
Eth1/2    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R
Eth1/3    Disabled  Disabled  PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R, MED-NWP
Eth1/4    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R, MED-CAP, MED-NWP
Eth1/5    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R
Eth1/6    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R
```

```
...
switch (config) # show lldp interface ethernet 1/4
TLV flags:
PD: port-description, SN: sys-name, SD: sys-description, SC: sys-capabilities, MA: management-address
ETS-C: ETS-Configuration, ETS-R: ETS-Recommendation, AP: Application Priority, PFC: Priority Flow Control
CEE: Converged Enhanced Ethernet DCBX version
MED-CAP: Media Capabilities
MED-NWP: MED-Network Policy

Interface Receive   Transmit  TLVs
-------------------------------------------------------------------------------------------
Eth1/4    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R, MED-CAP, MED-NWP.


switch (config) # show lldp interface ethernet 1/4 med-cap
Media Capabilities:
    LLDP-MED Capab   : Yes
    Network Policy   : Yes
    Location Id      : No
    Ext Power MDI-PSE: No
    Ext Power MDI-PD : No

Network Policy:
    Application Type : 1 (Voice)
    VLAN Id          : 11
    L2 Priority      : 0
    DSCP Value       : 0
```

To configure voice VLAN, take the following steps:

1. Create a VLAN.

```
switch (config) # vlan 200
switch (config vlan 200) # exit
switch (config) #
```

2. Set the interface mode to be hybrid.

```
switch (config) # interface ethernet 1/4 switchport mode hybrid
switch (config) # interface ethernet 1/4 switchport hybrid allowed-vlan 200
```

3. Assign the VLAN to the interface.

```
switch (config) # interface ethernet 1/4 switchport voice vlan 200
```

4. (Optional) Change the PVID of the port so that untagged packets go to a different VLAN than the default.

```
switch (config)# vlan 300
switch (config vlan 300)# exit
switch (config)# interface ethernet 1/4 switchport access vlan 300
```

5. Verify the configuration.

```
switch (config)# show interface switchport
Interface       Mode         Access vlan        Allowed vlans
-------------------------------------------------------------------------------
Eth1/1          access       1
Eth1/2          access       1
Eth1/3          access       1
Eth1/4          hybrid       300                200
Eth1/5          access       1
...
switch (config)# show lldp interface ethernet 1/4
TLV flags:
PD: port-description, SN: sys-name, SD: sys-description, SC: sys-capabilities, MA: management-address
ETS-C: ETS-Configuration, ETS-R: ETS-Recommendation, AP: Application Priority, PFC: Priority Flow Control
CEE: Converged Enhanced Ethernet DCBX version
MED-CAP: Media Capabilities
MED-NWP: MED-Network Policy

Interface Receive   Transmit  TLVs
-------------------------------------------------------------------------------------------
Eth1/4    Enabled   Enabled   PD, SN, SD, SC, MA, PFC, AP, ETS-C, ETS-R, MED-CAP, MED-NWP
switch (config)# show lldp interface ethernet 1/4 med-cap
Media Capabilities:
    LLDP-MED Capab   : Yes
    Network Policy   : Yes
    Location Id      : No
    Ext Power MDI-PSE: No
    Ext Power MDI-PD : No
```

```
Network Policy:
    Application Type : 1 (Voice)
    VLAN Id          : 200
    L2 Priority      : 0
    DSCP Value       : 0
```

To remove voice VLAN and LLDP-MED TLV, take the following steps:

1. Remove the voice VLAN from the interface.

```
switch (config)# no interface ethernet 1/4 switchport voice vlan
```

2. Disable the MED TLV from the interface.

```
switch (config)# interface ethernet 1/4 lldp med-tlv-select none
```

12.6.2 Limitations

1. LLDP MED cannot be enabled on a router port interface and vice versa (i.e. a port that has LLDP MED enabled cannot be configured as a router port interface).
2. LLDP MED cannot be enabled on a LAG and vice versa (i.e. a port that has LLDP MED enabled cannot be configured as a LAG).
3. If switchport is in trunk, dot1q-tunnel, or dcbx-access, configuring either the TLV or Voice VLAN gives a warning message.

12.7 Spanning Tree Protocol

The operation of Rapid Spanning Tree Protocol (RSTP) provides for rapid recovery of connectivity following the failure of a bridge/bridge port or a LAN. The RSTP component avoids this delay by calculating an alternate root port, and immediately switching over to the alternate port if the root port becomes unavailable. Thus, using RSTP, the switch immediately brings the alternate port to forwarding state, without the delays caused by the listening and learning states. The RSTP component conforms to IEEE standard 802.1D 2004.

RSTP enhancements is a set of functions added to increase the volume of RSTP in NVIDIA switches. It adds a set of capabilities related to the behavior of ports in different segments of the network. For example: the required behavior of a port connected to a non-switch entity, such as host, is to converge quickly, while the required behavior of a port connected to a switch entity is to converge based on the RSTP parameters.

Additionally, it adds security issues on a port and switch basis, allowing the operator to determine the state and role of a port or the entire switch should an abnormal event occur. For example: If a port is configured to be root-guard, the operator will not allow it to become a root-port under any circumstances, regardless of any BPDU that will have been received on the port.

12.7.1 Port Priority and Cost

When two ports on a switch are part of a loop, the STP port priority and port path cost configuration determine which port on the switch is put in the forwarding state and which port is put in the blocking state.

To configure port priority use the following command:

```
switch (config interface ethernet <inf>)# spanning-tree port-priority <0-240>
```

To configure port path cost use the following command:

```
switch (config interface ethernet <inf>)# spanning-tree cost <1-200000000>
```

12.7.2 Port Type

Port type has the following configuration options:

- edge - is not assumed to be converged by the RSTP learning/forwarding mechanism. It converges to forwarding quickly.

  > It is recommended to configure the port type for all ports connected to hosts as edge ports.

- normal - is assumed to be connected to a switch, thus it tries to be converged by the RSTP learning/forwarding. However, if it does not receive any BPDUs, it is operationally moved to be edge.
- network - is assumed to be connected only to a switch or bridge.

Each of these configuration options is mutually exclusive.

Port type is configured using the command spanning-tree port type. It may be applied globally on the switch (Config) level, which configures all switch interfaces. Another option is to configure ports individually by entering the interface's configuration mode.

- Global configuration:

  ```
  switch (config)# spanning-tree port type {edge , normal , network} default
  ```

- Interface configuration:

  ```
  switch (config interface ethernet <inf>)# spanning-tree port type {edge , normal, network}
  ```

For more information about this feature and its potential applications, please refer to the following community post:

- How To Configure Switch Port Types with NVIDIA Onyx

12.7.3 BPDU Filter

Using BPDU filter prevents the CPU from sending/receiving BPDUs on specific ports.

BPDU filtering is configured per interface. When configured, the port does not send any BPDUs and drops all BPDUs that it receives. To configure BPDU filter, use the following command:

```
switch (config interface ethernet <inf>)# spanning-tree bpdufilter {enable | disable}
```

12.7.4 BPDU Guard

BPDU guard is a security feature which, when enabled, will move the port to "down (suspended)" mode in case it receives BPDU packets. This feature becomes useful when connecting to an unauthorized switch.

To configure BPDU guard use the following command:

```
switch (config interface ethernet <inf>)# spanning-tree bpduguard {enable , disable}
```

12.7.4.1 Logging Example In Case of a BPDU Guard Event

```
Oct 29 22:55:30 r-anaconda-01 issd[7375]: TID
140652362820224: [issd.WARNING]: NPAPI_WRN: warning RstHandleInBpdu Received
BPDU on Port Eth1/12 with BPDU guard enabled. Disabling Port.
```

12.7.5 Loop Guard

Loop guard is a feature that prevents loops in the network.

When a blocking port in a redundant topology transitions to the forwarding state (accidentally), an STP loop occurs. This happens when BPDUs are no longer received by one of the ports in a physically redundant topology.

Loop guard is useful in switched networks where devices are connected point-to-point. A designated bridge cannot disappear unless it sends an inferior BPDU or brings the link down on a point-to-point connection.

> The loop guard configuration is only allowed on “network” and “normal” port types.

If loop guard is enabled and the port does not receive BPDUs, the port is put into an inconsistent state (blocking) until the port starts to receive BPDUs again. A port in the inconsistent state does not transmit BPDUs. If BPDUs are received again, loop guard alters its inconsistent state condition. STP converges to a stable topology without the failed link or bridge after loop guard isolates the failure.

Disabling loop guard moves all loop-inconsistent ports to listening state.

To configure loop guard use the following command:

```
switch (config interface ethernet <inf>)# spanning-tree guard loop
```

12.7.6 Root Guard

Configuring root guard on a port prevents that port from becoming a root port. A port put in root-inconsistent (blocked) state if an STP convergence is triggered by a BPDU that makes that port a root port. The port is unblocked after the port stops sending BPDUs.

To configure loop guard use the following command:

```
switch (config interface ethernet <inf>)# spanning-tree guard root
```

12.7.7 MSTP

Spanning Tree Protocol (STP) is a mandatory protocol to run on L2 Ethernet networks to eliminate network loops and the resulting broadcast storm caused by these loops. Multiple STP (MSTP) enables the virtualization of the L2 domain into several VLANs, each governed by a separate instance of a spanning tree which results in a network with higher utilization of physical links while still keeping the loop free topology on a logical level.

Up to 64 MSTP instances can be configured on a switch.

12.7.8 RPVST

Rapid Per-VLAN Spanning Tree (RPVST) flavor of the STP provides finer-grained traffic by paving a spanning-tree instance per each configured VLAN. Like MSTP, it allows a better utilization of the network links comparing to RSTP.

The following figure exhibits a typical RPVST network configuration to get a better utilization on the inter-switch trunk ports.

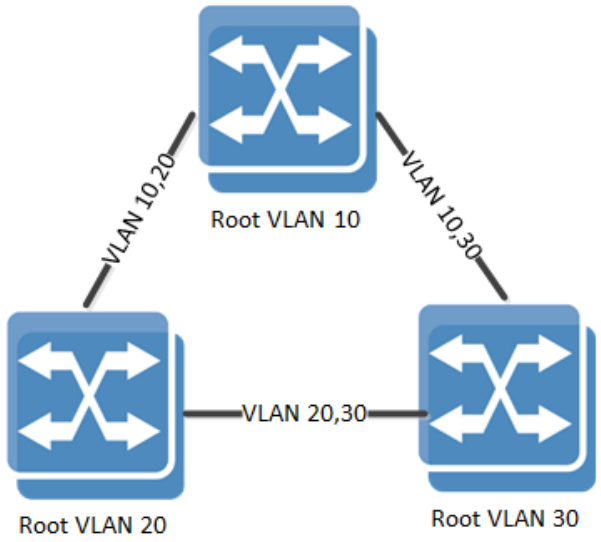

12.7.8.1 RPVST and VLAN Limitations

When the STP of the switch is set to RPVST, spanning tree is set on each of the configured VLANs in the system by default. To enable the spanning tree mode, the command “spanning-tree” must be run.

Each VLAN runs an STP state machine and an RPVST instance. There is a global limitation on the number of active state machines that can operate inNVIDIA Onyx. Enforcement of this limitation is done through the maximum number of VLANs allowed in the system (128).

The state machine takes attributes like forward time, hello time, max age and priority, etc.

> When configuring priority on a VLAN in RPVST, the operational priority given to the VLAN is a summation of what the user configured and the value of the VLAN itself. For example, running “spanning-tree vlan 10 priority 32768” yields a priority of 32778 for VLAN 10.

12.7.8.2 RPVST and RSTP Interoperability

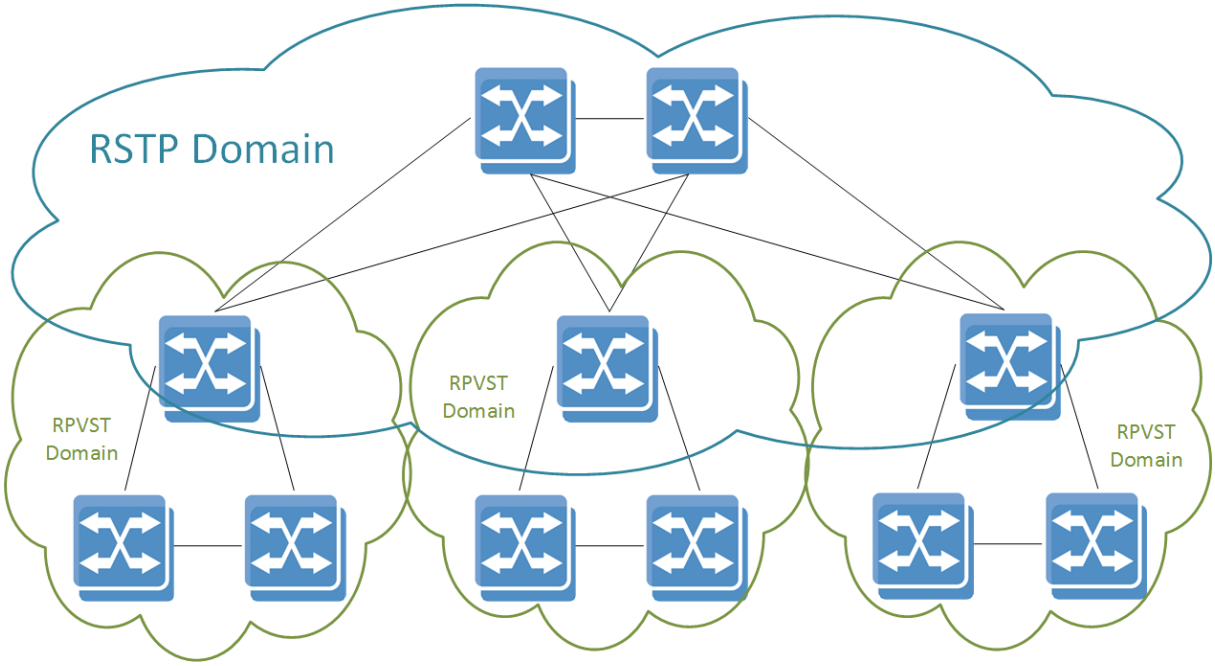

RPVST domains can be interconnected by a standard 802.1Q domain that runs RSTP protocol. While the RSTP domain builds a single common instance spanning tree, the RPVST domains at the edge continue to build a tree per VLAN while exchanging tagged RPVST multicast BPDUs.
(This exchange may happen on untagged RPVST BPDUs as well.) The switch devices that are in the boundary between the RPVST and the RSTP domains should be configured as RPVST mode.

When set to RPVST mode, the switch continues to run the common instance spanning tree (CIST) state machine on VLAN 1 by exchanging IEEE BPDUs with the legacy RSTP switches.

To successfully connect RSTP and RPVST domains, the system administrator must align the native VLAN configuration across all network switches, or in other words, the internal identification of untagged packets to VLAN.

12.7.9 STP Commands

12.7.9.1 spanning-tree

	spanning-tree no spanning-tree Globally enables spanning tree. The no form disables spanning tree.
Syntax Description	N/A
Default	Spanning tree is enabled
Configuration Mode	config
History	3.1.0000
Example	`switch (config) # no spanning-tree`
Related Commands	show spanning-tree
Notes	

12.7.9.2 spanning-tree mode

	spanning-tree mode {mst | rst | rpvst} no spanning-tree mode Changes spanning tree mode. The no form of the command sets the parameter to its default value.	
Syntax Description	mst	Multiple spanning tree
	rst	Rapid spanning tree
	rpvst	Rapid per-VLAN spanning tree
Default	rst	
Configuration Mode	config	
History	3.3.4150	
Example		
Related Commands	show spanning-tree	
Notes	The number of VLANs supported by RPVST is 128	

12.7.9.3 spanning-tree (timers)

	spanning-tree [forward-time <time in secs> | hello-time <time in secs> | max-age <time in secs>] no spanning-tree [forward-time | hello-time | max-age | priority] Configures spanning tree timers. The no form of the command sets the timer to default.	
Syntax Description	forward-time	Controls how fast a port changes its spanning tree state from Blocking state to Forwarding state Parameter range: 4-30 seconds

	hello-time	Determines how often the switch broadcasts its hello message to other switches when it is the root of the spanning tree Parameter range: 1-2 seconds
	max-age	Sets the maximum age allowed for the Spanning Tree Protocol information learnt from the network on any port before it is discarded Parameter range: 6-40 seconds
Default	forward-time: 15 seconds hello-time:2 seconds max-age: 20 seconds	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # spanning-tree forward-time`	
Related Commands	show spanning-tree	
Notes	The following formula applies on the spanning tree timers: 2*(ForwardTime -1)>=MaxAgeTime >= 2*(Hello Time + 1)	

12.7.9.4 spanning-tree port type (default global)

	spanning-tree port type {edge [bpdufilter | bpduguard] | network [bpduguard] | normal [bpduguard]} default no spanning-tree port type default Configures all switch interfaces as edge/network/normal ports. These ports can be connected to any type of device. The no form of the command disables the spanning tree operation.	
Syntax Description	edge	Assumes all ports are connected to hosts/servers
	bpdufilter	Configures to enable the spanning tree BPDU filter
	bpduguard	Configures to enable the spanning tree BPDU guard
	network	Assumes all ports are connected to switches and bridges
	normal	The port type (edge or network) determines according to the spanning tree operational mode
Default	normal	
Configuration Mode	config	
History	3.1.0000	
	3.4.0008	Updated command syntax
Example	`switch (config) # spanning-tree port type edge default`	
Related Commands	show spanning-tree	
Notes		

12.7.9.5 spanning-tree priority

	spanning-tree priority <bridge-priority> no spanning-tree priority Sets the spanning tree bridge priority. The no form of the command sets the bridge priority to default.	
Syntax Description	bridge-priority	Sets the bridge priority for the spanning tree Value must be in increments of 4096, starting from 0 (accepted values: 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, 61440)
Default	32786	
Configuration Mode	config	
History	3.1.0000	
Example	`switch (config) # spanning-tree priority 4096`	
Related Commands	show spanning-tree	
Notes		

12.7.9.6 spanning-tree port-priority

	spanning-tree port-priority <priority> no spanning-tree port-priority Configures the spanning-tree interface priority. The no form of the command returns configuration to its default.	
Syntax Description	priority	Spanning tree interface priority Possible values: 0, 16, 32,48, 64, 80, 96, 112, 128,144, 160, 176, 192, 208, 224, 240
Default	128	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.3.4500	Added MPO configuration mode
Example	`switch (config interface ethernet 1/1) # spanning-tree port-priority 16`	
Related Commands	show spanning-tree	
Notes		

12.7.9.7 spanning-tree cost

	spanning-tree cost <port cost> no spanning-tree cost Configures the interface cost of the spanning tree. The no form of the command returns configuration to its default.	
Syntax Description	port cost	Sets the spanning tree cost of an interface. Range: 0-200000000

Default	The default cost is derived from the interface speed: • 1Gb/s 20000 • 10Gb/s 2000 • 40Gb/s 500 • 50Gb/s 400 • 100Gb/s 200	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.3.4500	Added MPO configuration mode
Example	`switch (config interface ethernet 1/1) # spanning-tree cost 1000`	
Related Commands	show spanning-tree	
Notes	• LAG default cost is calculated by dividing the port speed by the number of active links in UP state. For example: if there were 4 links in the LAG out of which only two are in UP state, assuming the port speed is 10Gbps, the LAG cost will be 2000/2 = 1000. • When configuring the cost for a LAG, the cost will be fixed to this configuration, no matter what the number of active links (UP state) in the LAG is • Unstable network may cause the LAG cost to change dynamically assuming the cost parameter is not configured for anything else other than default	

12.7.9.8 spanning-tree port type

	spanning-tree port type <port type> no spanning-tree port type Configures spanning-tree port type The no form of the command returns configuration to default.	
Syntax Description	default	According to global configuration.
	edge	Assumes all ports are connected to hosts/servers.
	normal	The port type (edge or network) determines according to the spanning tree operational mode.
	network	Assumes all ports are connected to switches and bridges.
	bpdufilter	Configures to enable the spanning tree BPDU filter.
	bpduguard	Configures to enable the spanning tree BPDU guard.
Default	Globally defined by the command “spanning-tree port type <port-type> default”.	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.3.4500	Added MPO configuration mode
Example	`switch (config interface ethernet 1/1) # spanning-tree port type edge`	
Related Commands	show spanning-tree	
Notes		

12.7.9.9 spanning-tree guard

	spanning-tree guard {loop | root} no spanning-tree guard {loop | root} Configures spanning-tree guard. The no form of the command returns configuration to default.	
Syntax Description	loop	Enables loop-guard on the interface. If the loop-guard is enabled, upon a situation where the interface fails to receive BPDUs the switch will not egress data traffic on this interface.
	root	Enables root-guard on the interface. If root-guard is enabled on the interface, the interface will never be selected as root port.
Default	loop-guard and root-guard are disabled	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
	3.3.4500	Added MPO configuration mode
Example	`switch (config interface ethernet 1/1) # spanning-tree guard root`	
Related Commands	show spanning-tree	
Notes		

12.7.9.10 spanning-tree bpdufilter

	spanning-tree bpdufilter {disable | enable} no spanning-tree bpdufilter Configures spanning-tree BPDU filter on the interface. The interface will ignore any BPDU that it receives and will not send PDBUs, The STP state on the port will move to the forwarding state. The no form of the command returns the configuration to default.	
Syntax Description	disable	Disables the BPDU filter on this port
	enable	Enables the BPDU filter on this port
Default	BPDU filter is disabled	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.0000	
Example	`switch (config interface ethernet 1/1) # spanning-tree bpdufilter enable`	
Related Commands	show spanning-tree	
Notes	This command can be used when the switch is connected to hosts	

12.7.9.11 clear spanning-tree counters

	clear spanning-tree counters Clears the spanning-tree counters.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Example	`switch (config) # clear panning-tree counters`
Related Commands	show spanning-tree
Notes	

12.7.9.12 spanning-tree mst max-hops

	spanning-tree mst max-hops <max-hops> no spanning-tree mst max-hops Specifies the max hop value inserts into BPDUs that sent out as the root bridge. The no form of the command sets the parameter to its default value.	
Syntax Description	max-hops	Max hop value Range: 6-40
Default	20	
Configuration Mode	config	
History	3.3.4150	
Example	`switch (config) # spanning-tree mst max-hops 20`	
Related Commands	show spanning-tree	
Notes	• The max hop setting determines the number of bridges in an MST region that a BPDU can traverse before it is discarded • This command is available when global STP mode is set to MST	

12.7.9.13 spanning-tree mst priority

	spanning-tree mst <mst-instance> priority <priority> no spanning-tree mst <mst-instance> priority Configures the specified instance's priority number. The no form of the command sets the parameter to its default value.	
Syntax Description	mst-instance	MST instance Range: 1-64
	priority	MST instance port priority Value must be in increments of 4096, starting from 0 (accepted values: 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, 61440)
Default	32768	
Configuration Mode	config	

History	3.3.4150
Example	`switch (config) # spanning-tree mst 1 priority 32768`
Related Commands	show spanning-tree
Notes	• The bridge priority is the four most significant digits of the bridge ID, which is used by spanning tree algorithms to select the root bridge and choose among redundant links. Bridge ID numbers range from 0-65535 (16 bits); bridges with smaller bridge IDs are elected over other bridges. • This command is available when global STP mode is set to MST

12.7.9.14 spanning-tree mst vlan

	spanning-tree mst <mst-instance> vlan <vlan-id> no spanning-tree mst <mst-instance> vlan <vlan-id> Maps a VLAN or a range of VLANs into an MSTP instance. The no form of the command unmaps a VLAN or a range of VLANs from MSTP instances.	
Syntax Description	mst-instance	MST instance Range: 1-64
	vlan-id	A single VLAN or a a range of VLANs Formats: “<vlan>” or “<from-vlan>-<to-vlan>” (see Example below)
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
Example	`switch (config) # spanning-tree mst 1 vlan 10-20`	
Related Commands	show spanning-tree	
Notes	This command is available when global STP mode is set to MST	

12.7.9.15 spanning-tree mst revision

	spanning-tree mst revision <number> no spanning-tree mst revision Configures the MSTP revision number. The no form of the command sets the parameter to its default value.	
Syntax Description	number	MST revision number Range: 0-65535
Default	0	
Configuration Mode	config	
History	3.3.4150	
Example	`switch (config)# spanning-tree mst revision 1`	
Related Commands	show spanning-tree	
Notes	• The revision number is one of three parameters, along with the MST name and VLAN-to-instance map, that identify the switch’s MST region • This command is available when global STP mode is set to MST	

12.7.9.16 spanning-tree mst name

	spanning-tree mst name <name> no spanning-tree mst name Configures the MSTP name. The no form of the command sets the parameter to its default value.	
Syntax Description	name	MST name: Up to 32 characters
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
Example	`switch (config)# spanning-tree mst name mymst`	
Related Commands	show spanning-tree	
Notes	• The name is one of three parameters, along with the MST revision number and VLAN-to-instance map, that identifies the switch's MST region • This command is available when global STP mode is set to MST	

12.7.9.17 spanning-tree mst root

	spanning-tree mst <mst-instance> root <role> no spanning-tree mst <mst-instance> root Changes the bridge priority for the specified MST instance to the following values: • Primary – 8192 • Secondary – 16384 The no form of the command sets the parameter to its default value.	
Syntax Description	mst-instance	MSTP instance Range: 1-64
	role	Possible values: "primary" or "secondary"
Default	primary	
Configuration Mode	config	
History	3.3.4150	
	3.7.1000	Updated example
Example	`switch (config)# spanning-tree mst 1 root primary`	
Related Commands	show spanning-tree	
Notes	• The root command is a way to automate a system configuration while 'playing' with the priority field. The priority field granularity may be too explicit for some users in case you wish to have 2 levels of priority (primary and secondary). So by default all the switches get the same priority and while using the root option you can get the role of master and backup by setting the priority field to a predefined value. • This command is available when global STP mode is set to MST	

12.7.9.18 spanning-tree mst port-priority

	spanning-tree mst <mst-instance> port-priority <priority> no spanning-tree mode Changes the spanning tree mode. The no form of the command sets the parameter to its default value.	
Syntax Description	mst-instance	MST instance Range: 1-64
	priority	MST instance port priority Valid values are: 0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 192, 208, 224 and 240
Default	rst	
Configuration Mode	config interface ethernet config interface port-channel	
History	3.3.4150	
Example	`switch (config interface ethernet 1/1)# spanning-tree mst 1 port-priority 32768`	
Related Commands	show spanning-tree	
Notes	This command is available when global STP mode is set to MST	

12.7.9.19 spanning-tree mst cost

	spanning-tree mst <mst-instance> cost <cost-value> no spanning-tree mode Configures the cost per MSTP instance. The no form of the command sets the parameter to its default value.	
Syntax Description	mst-instance	MST instance Range: 1-64
	cost-value	MST instance port cost Range: 0-200000000
Default	• 20000 for 1Gb/s • 2000 for 10Gb/s • 500 for 40Gb/s • 357 for 56Gb/s • 200 for 100Gb/s	
Configuration Mode	config interface ethernet config interface port-channel	
History	3.3.4150	
Example	`switch (config interface ethernet 1/1)# spanning-tree mst 1 cost 4000`	
Related Commands	show spanning-tree	
Notes	This command is available when global STP mode is set to MST	

12.7.9.20 spanning-tree vlan forward-time

	spanning-tree vlan <vid> forward-time <secs> no spanning-tree vlan <vid> forward-time Configures how fast an interface changes its spanning tree state from Blocking to Forwarding. The no form of the command resets the parameter value to its default.	
Syntax Description	secs	Parameter range: 4-30 seconds.
Default	15 seconds	
Configuration Mode	config	
History	3.4.1100	
Example	`switch (config) # spanning-tree vlan 10 forward-time 15`	
Related Commands	show spanning-tree	
Notes	• The following formula applies on the spanning tree timers: 2*(ForwardTime -1)>=MaxAgeTime >= 2*(Hello Time + 1) • This command is available when global STP mode is set to RPVST	

12.7.9.21 spanning-tree vlan hello-time

	spanning-tree vlan <vid> hello-time <secs> no spanning-tree vlan <vid> hello-time Configures how often the switch broadcasts its hello message to other switches when it is the root of the spanning tree. The no form of the command resets the parameter value to its default.	
Syntax Description	vid	VLAN ID
	secs	Range: 1-2 seconds
Default	2 seconds	
Configuration Mode	config	
History	3.4.1100	
Example	`switch (config) # spanning-tree vlan 10 hello-time 2`	
Related Commands	show spanning-tree	
Notes	• The following formula applies on the spanning tree timers: 2*(ForwardTime -1)>=MaxAgeTime >= 2*(Hello Time + 1) • This command is available when global STP mode is set to RPVST	

12.7.9.22 spanning-tree vlan max-age

	spanning-tree vlan <vid> max-age <secs> no spanning-tree vlan <vid> max-age Sets the maximum age allowed for the Spanning Tree Protocol information learned from the network on any port before it is discarded. The no form of the command resets the parameter value to its default.	
Syntax Description	secs	Range: 6-40 seconds
Default	20 seconds	

Configuration Mode	config
History	3.4.1100
Example	`switch (config) # spanning-tree vlan 10 max-age 20`
Related Commands	show spanning-tree
Notes	

12.7.9.23 spanning-tree vlan priority

	spanning-tree vlan <vid> priority <priority> no spanning-tree vlan <vid> priority Configures RPVST instance port priority. The no form of the command resets the parameter value to its default.	
Syntax Description	vid	VLAN ID
	priority	MST instance port priority Value must be in increments of 4096, starting from 0 (accepted values: 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, 61440)
Default	32768	
Configuration Mode	config	
History	3.4.1100	
Example	`switch (config) # spanning-tree vlan 10 priority 32768`	
Related Commands	show spanning-tree	
Notes		

12.7.9.24 show spanning-tree

	show spanning-tree Displays spanning tree information.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.4.1100	Updated example with R and G flags
	3.6.6000	Updated example
	3.6.6102	Added note on MLAG spanning-tree cost
Example		

```
switch (config) # show spanning-tree

Switch                      : ethernet-default
Spanning tree protocol rst : enabled
Spanning tree force version: 2

Root ID:
  Priority: 32768
  Address : 7c:fe:90:ff:2c:40

  This bridge is the root

  Hello Time (sec)    : 2
  Max Age (sec)       : 20
  Forward Delay (sec): 15

Bridge ID:
  Priority            : 32768
  Address             : 7c:fe:90:ff:2c:40
  Hello Time (sec)    : 2
  Max Age (sec)       : 20
  Forward Delay (sec): 15

L: Loop Inconsistent
R: Root Inconsistent
G: BPDU Guard Inconsistent

------------------------------------------------------------------------
Interface         Role         Sts              Cost      Prio    Type
------------------------------------------------------------------------
Eth1/7            Designated   Discarding       200       128     normal
Eth1/8            Disabled     Discarding(G)    200       128     edge
```

Related Commands	clear spanning-tree counters spanning-tree
Notes	• MLAG spanning-tree cost is always equal to the cost of there being 2 member ports in the MLAG (even if one of the member ports fails or a new port is added) • If a port is in BPDU Guard inconsistent mode, the interface status will move to "down (suspended)".

12.7.9.25 show spanning-tree detail

	show spanning-tree detail Displays detailed spanning-tree configuration and statistics.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.6.4110	Updated example
	3.6.5000	Updated example
Example		
Related Commands	clear spanning-tree counters spanning-tree	
Notes		

12.7.9.26 show spanning-tree interface

	show spanning-tree interface {ethernet <slot>/<port> | port-channel <port-channel> | mlag-port-channel <mlag-port-channel> Display running state for specific interfaces.	
Syntax Description	ethernet	Ethernet interface
	port-channel	LAG instance
	mlag-port-channel	MLAG instance
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4150	

Example	`switch (config) # show spanning-tree  1/2`
	`Eth1/2 is Disabled  Discarding`
	`Port path cost 500, Port priority 128, Port Identifier 128.5`
	`Designated root has priority 0, address unknown`
	`Designated bridge has priority 0, address unknown`
	`Designated port id 0.0, designated path cost 0`
	`Number of transitions to forwarding state: 0`
	`Port type: normal`
	`PortFast is: off`
	`Bpdu filter: disabled`
	`Bpdu guard: disabled`
	`Loop guard: disabled`
	`Root guard: disabled`
	`Link type: point-to-point`
	`BPDU: sent: 0   received: 0`	

Related Commands	clear spanning-tree counters spanning-tree	
Notes		

12.7.9.27 show spanning-tree mst

	show spanning-tree mst [details | <instance> interface {ethernet <slot>/<port> | port-channel <port-channel> | mlag-port-channel <mlag-port-channel>}] Displays basic multi-spanning-tree information.	
Syntax Description	details	Displays detailed multi-spanning-tree configuration and statistics
	ethernet	Ethernet interface
	port-channel	LAG instance
	mlag-port-channel	MLAG instance
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4150	
	3.6.6000	Updated example
Example		

```
switch (config) # switch (config) # show spanning-tree mst

MST0:
  vlans mapped: 1-1023,1025-2047,2049-3071,3073-4094

L: Loop Inconsistent
R: Root Inconsistent
G: BPDU Guard Inconsistent

--------------------------------------------------------------------------
Interface         Role         Sts              Cost      Prio   Type
--------------------------------------------------------------------------
Eth1/7            Designated   Discarding       200       128.7  normal
Eth1/8            Disabled     Discarding(G)    200       128.8  edge
```

Related Commands	clear spanning-tree counters spanning-tree
Notes	

12.7.9.28 show spanning-tree root

	show spanning-tree root Displays root multi-spanning-tree information.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.4150
Example	

```
switch (config) # show spanning-tree root

Instance   Priority   MAC addr           Root Cost  Hello Time  Max Age    FWD Dly    Root Port
-------    ------     --------           ---------  ----------  --------   -------    ---------
MST0       32768      00:02:c9:71:ed:40  500        2           20         15         Eth1/20
MST1       32768      00:02:c9:71:f0:c0  0          2           20         15         -
MST2       0          00:02:c9:71:f0:c0  0          2           20         15         -
MST3       32768      00:02:c9:71:f0:c0  0          2           20         15         -
```

Related Commands	clear spanning-tree counters spanning-tree
Notes	

12.7.9.29 show spanning-tree vlan

	show spanning-tree vlan <vid> [detail | interface {ethernet <slot>/<port> | port-channel <port-channel> | mlag-port-channel <mlag-port-channel>}] Displays spanning-tree protocol information.	
Syntax Description	vid	VLAN ID. Range is also supported Format: <vid1>[-<vid2>]
	detail	Displays detailed RPVST configuration and statistics
	ethernet	Ethernet interface
	port-channel	LAG instance
	mlag-port-channel	MLAG instance

Default	N/A	
Configuration Mode	Any command mode	
History	3.4.1100	
	3.6.5000	Updated example output

Example

```
switch (config) # show spanning-tree vlan 1 detail

Switch ethernet-default
Spanning tree protocol is enabled
Bridge is executing the rpvst compatible Spanning Tree Protocol

Vlan 1:
    Bridge Identifier priority: 32769
    Bridge Identifier address: e4:1d:2d:3d:5e:c0
    Configured hello time: 2, max age 20, forward delay 15
    Current root: priority 32769, address e4:1d:2d:3d:5e:c0
    Number of topology changes: 0, last change occurred 00:00:00 ago
    Last TCN received from: N/A
    Timers: hold 6 hello 2, max age 20, forward delay 15
    Default port type: normal
    Default bpdu filter: disabled
    Default bpdu guard: disabled
```

Related Commands	clear spanning-tree counters spanning-tree
Notes	

12.7.9.30 show spanning-tree vlan topo-change-history

	show spanning-tree vlan <vid> topo-change-history Displays spanning-tree topology change notification history per VLAN.	
Syntax Description	vid	VLAN ID Format: <vid1>[-<vid2>]
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4110	

Example

```
switch (config) # show spanning-tree vlan 50 topo-change-history

Vlan 50

------------------------------------
Interface    Date          Time
------------------------------------
Eth1/49      07/18/17      04:39:58
Eth1/49      07/18/17      04:39:55
Eth1/49      07/18/17      04:38:11
Eth1/49      07/18/17      04:38:09
```

Related Commands	spanning-tree
Notes	

12.7.9.31 show spanning-tree mst topo-change-history

	show spanning-tree mst <mst-instance> topo-change-history Displays spanning-tree topology change notification history per instance.

Syntax Description	mst-instance	MST instance Range: 1-64
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4110	

Example	`switch (config) # show spanning-tree mst 5 topo-change-history`
	
	`Instance 5`
	
	`------------------------------------`
	`Interface  Date  Time`
	`------------------------------------`
	`Eth1/49  07/18/17  04:43:51`
	`Eth1/49  07/18/17  04:43:33`	

Related Commands	spanning-tree	
Notes		

12.7.9.32 show spanning-tree topo-change-history

	show spanning-tree topo-change-history Displays spanning-tree topology change notification history.	
Syntax Description	mst-instance	MST instance Range: 1-64
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4110	

Example	`switch (config) # show spanning-tree topo-change-history`
	
	`------------------------------------`
	`Interface  Date  Time`
	`------------------------------------`
	`Eth1/49  07/27/17  09:39:38`
	`Eth1/35  07/27/17  09:35:42`
	`Eth1/35  07/27/17  09:35:40`
	`Eth1/35  07/27/17  09:35:08`
	`Eth1/35  07/27/17  09:35:06`
	`Eth1/35  07/27/17  09:32:05`
	`Eth1/35  07/27/17  09:32:03`
	`Eth1/35  07/27/17  09:31:42`
	`Eth1/35  07/27/17  09:31:40`	

Related Commands	spanning-tree	
Notes		

12.8 MAC Address Table

12.8.1 Configuring Unicast Static MAC Address

You can configure static MAC addresses for unicast traffic. This feature improves security and reduces unknown unicast flooding.

To configure Unicast Static MAC address, run the following:

```
mac-address-table static unicast <destination mac address> vlan <vlan identifier(1-4094)> interface ethernet
<slot>/<port>
```

For example:

```
switch (config) # mac-address-table static 00:11:22:33:44:55 vlan 1 interface ethernet 1/1
```

12.8.2 MAC Learning Considerations

MAC learning may be disabled using the command mac-learning disable which is beneficial in the following situations:

- To prevent denial-of-service attacks
- To manage the available MAC address table space by controlling which interfaces can learn MAC addresses
- To duplicate to a dedicated server (port7 in the figure below) all the packets that one host (host1; port1) sends to another (host2; port2), like in port mirroring. To accomplish this, MAC learning is disabled on port2. In this case the FDB does not obtain the MAC address of host2. Also, to prevent broadcast to every port, it is possible to configure a VLAN (VLAN 80) which ports 1, 2 and 7 are member of.

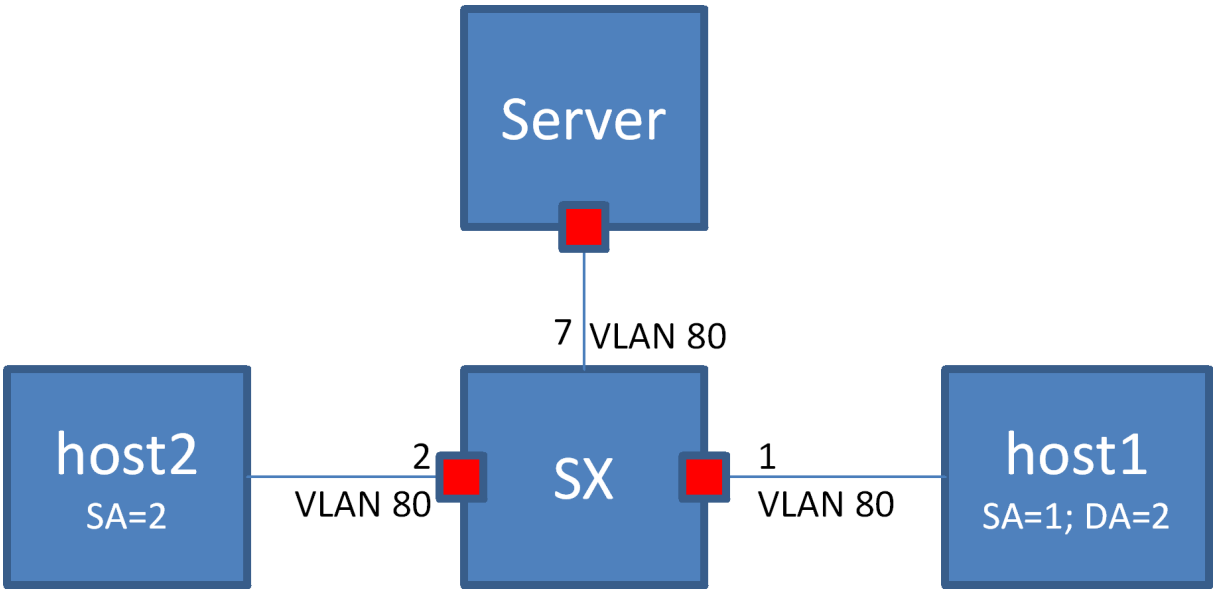

12.8.3 MAC Address Table Commands

12.8.3.1 mac-address-table aging-time

	mac-address-table aging-time <age> no mac-address-table aging-time Sets the maximum age of a dynamically learnt entry in the MAC address table. The no form of the command resets the aging time of the MAC address table to its default.	
Syntax Description	age	10-1000000 seconds
Default	300	
Configuration Mode	config	
History	3.1.0600	
Example	`switch (config) # mac-address-table aging-time 50`	
Related Commands	show mac-address-table show mac-address-table aging time	
Notes		

12.8.3.2 mac-address-table static

	mac-address-table static <mac address> vlan <vlan> interface <if-type> <if-number> no mac-address-table static <mac address> vlan <vlan> interface <if-type> <if-number> Configures a static MAC address in the forwarding database. The no form of the command deletes a configured static MAC address from the forwarding database.	
Syntax Description	mac address	Destination MAC address
	vlan	VLAN ID or VLAN range
	if-type	Ethernet or port-channel interface type
	if-number	Interface number (i.e. 1/1, 3)
Default	No static MAC addresses available in default	
Configuration Mode	config	
History	3.1.0600	
Example	`switch (config) # mac-address-table static aa:aa:aa:aa:aa:aa vlan 1 interface ethernet 1/7`	
Related Commands	show mac-address-table mac-address-table aging time	
Notes	The no form of the command will not clear a dynamic MAC address. Dynamic MAC addresses are cleared using the “clear mac-address-table dynamic” command.	

12.8.3.3 mac-learning disable

	mac-learning disable no mac-learning disable Disables MAC-address learning. The no form of the command enables MAC-address learning.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config interface ethernet config interface port-channel
History	3.1.0600
Example	`switch (config interface ethernet 1/1) # mac-learning disable`
Related Commands	
Notes	• When adding a port to a LAG, the port needs to be aligned with the LAG's configuration • When removing a port from a LAG, the port remains in whichever configuration the LAG is in • Disabling MAC learning is not supported on a local analyzer port. • Disabling MAC learning is not supported on an IPL LAG.

12.8.3.4 clear mac-address-table dynamic

	clear mac-address-table dynamic Clear the dynamic entries in the MAC address table.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0600
Example	`switch (config) # clear mac-address-table dynamic`
Related Commands	mac-address-table aging-time mac-address-table static show mac-address-table
Notes	This command does not clear the MAC addresses learned on the mgmt0 port. Static entries are deleted using the "no mac-address-table static" command.

12.8.3.5 show mac-address-table

	show mac-address-table [address <mac-address> | <if-number> | vlan [<vlan> | range <range>] | unicast] Displays the static and dynamic unicast and multicast MAC addresses for the switch. Various of filter options available.	
Syntax Description	mac-address	Filters the table to a specific MAC address.
	if-number	Filters the table to a specific interface.
	vlan	Filters the table to a specific VLAN number (1-4094).

	range	Filters the table to a range of VLANs.
	unicast	Filters the table to a unicast addresses only.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0600	
	3.3.4500	Updated example
	3.8.1000	Updated syntax & example
Example	switch (config) # show mac-address-table Switch ethernet-default Vlan Mac Address Type Interface ---- ----------- ---- ------------ 1 00:00:00:00:00:01 Static Po5 1 00:00:3d:5c:fe:16 Dynamic Eth1/1 1 00:00:3d:5d:fe:1b Dynamic Eth1/2 Number of unicast: 2 switch (config) # show mac-address-table unicast --- Vlan Mac Address Type Port\Next Hop --- 1 24:8a:07:2e:61:72 Dynamic Eth1/31 6 00:00:11:22:33:44 Static 192.168.2.2(nve1) 6 00:00:66:77:88:99 Static 192.168.2.2(nve1)	
Related Commands	mac-address-table static clear mac-address-table	
Notes		

12.8.3.6 show mac-address-table aging-time

	show mac-address-table aging-time Displays the MAC address table aging time.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0600
Example	switch (config) # show mac-address-table aging-time Mac Address Aging Time: 300
Related Commands	mac-address-table aging-time mac-address-table static clear mac-address-table
Notes	MAC addresses learned on the mgmt0 is not shown by this command.

12.8.3.7 show mac-address-table interface

	show mac-address-table interface [port-channel | mlag-port-channel <if>] Displays the MAC address table of a LAG or an MPO.
Syntax Description	N/A

Default	N/A
Configuration Mode	Any command mode
History	3.6.4006
Example	switch (config) # show mac-address-table ---------- Vlan Mac Address Type Port ---------- 1 E4:1D:2D:37:11:22 Static Eth1/1 1 E4:1D:2D:37:3E:11 Static Po5 Number of unicast: 2 Number of multicast: 0 switch (config) # show mac-address-table interface port-channel 5 ---------- Vlan Mac Address Type Port ---------- 1 E4:1D:2D:37:3E:11 Static Po5 Number of unicast: 1 Number of multicast: 0
Related Commands	mac-address-table static clear mac-address-table
Notes	

12.8.3.8 show mac-address-table interface nve

	show mac-address-table interface nve <nve-id> Displays MAC address table on specific NVE interface.	
Syntax Description	nve-id	NVE ID
Default	N/A	
Configuration Mode	Any command mode	
History	3.8.1000	
Example	switch (config) # show mac-address-table interface nve 1 ---------- Vlan Mac Address Type Port\Next Hop ---------- 60 E4:1D:2D:37:11:22 Dynamic Number of unicast(local): 1 Number of NVE: 1	
Related Commands	protocol nve mac-address-table static clear mac-address-table	
Notes	This command is not supported if NVE is not enabled.	

12.8.3.9 show mac-address-table summary

	show mac-address-table summary Displays total number of unicast/multicast MAC address entries.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode

History	3.6.2002	
	3.8.1000	Updated example

Example	`switch (config) # show mac-address-table summary`
	`Number of unicast(local): 4`
	`Number of NVE: 2`	

Related Commands	mac-address-table static clear mac-address-table	
Notes		

12.9 MLAG

A link aggregation group (LAG) is used for extending the bandwidth from a single link to multiple links and provide redundancy in case of link failure. Extending the implementation of the LAG to more than a single device provides yet another level of redundancy that extends from the link level to the node level. This extrapolation of the LAG from single to multiple switches is referred to as multi-chassis link aggregation (MLAG). MLAG is supported on Ethernet blades' internal and external ports.

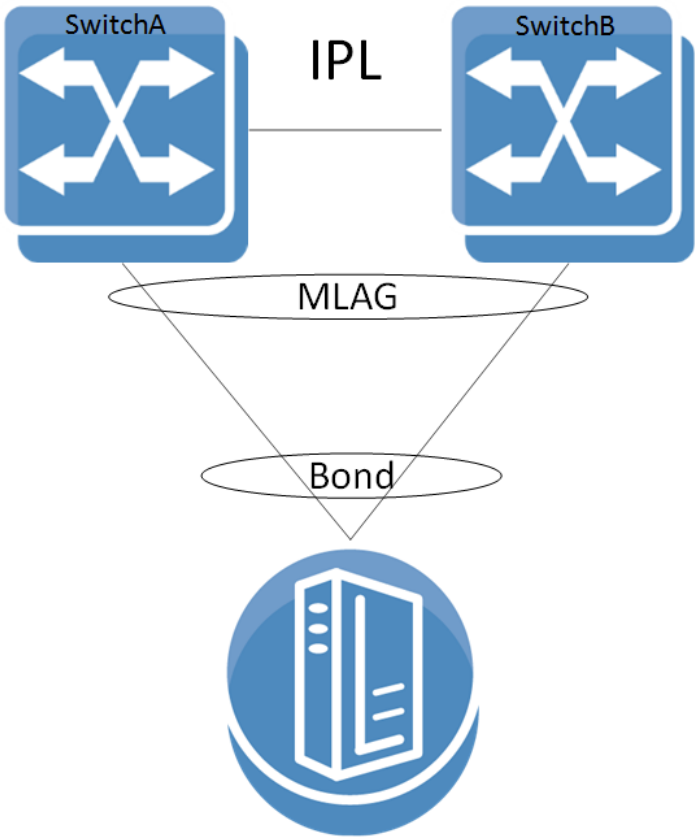

Each switch configuration is independent and the user is responsible for configuring both switches similarly pertaining MLAG (e.g., MLAG port-channel VLAN membership, static MAC, ACL, and so forth).

A peered device (host or switch) connecting to switches running an MLAG runs a standard LAG and is unaware of the fact that the LAG connects to two separate switches.

The MLAG switches share an inter-peer link (IPL) between them for carrying control messages in a steady state or data packages in failure scenarios. Thus, the bandwidth of the IPL should be defined accordingly. The IPL itself can be a LAG and may be constructed of links of any supported speed. In such a case, PFC must be configured on this IPL. The figure in section "Configuring MLAG" illustrates this. The IPL serves the following purposes:

- MLAG protocol control: keepalive messages, MAC sync, MLAG port sync, and so forth
- MLAG port failure: serves redundancy in case of a fallen link on one of the MLAG switches
- Layer-3 failure: serves redundancy in case of a failed connection between the MLAG switches and the rest of the L3 network should there be one

> The IPL VLAN interface must be used only for MLAG protocol and must not be used by any other interfaces (e.g., LAG, Ethernet).

> Ports 21 and 22 are dedicated IPL ports for MLAG protocol on the SH2200 switch system.

The MLAG protocol is made up of the following components to be expanded later:

- Keepalive
- Unicast and multicast sync
- MLAG port sync

When positioned at the top of rack (ToR) and connecting with a Layer-3 uplink, the MLAG pair acts as the L3 border for the hosts connected to it. To allow default gateway redundancy, both MLAG switches should be addressed by the host via the same default gateway address.

MLAG uses an IP address (VIP) that points to all MLAG member nodes.

When running MLAG as L2/L3 border point, an MAGP VIP must be deployed as the default gateway for MLAG port-channels (MPOs).

> When MLAG is connected through a Layer-2 based uplink, there is no need to apply default gateway redundancy towards hosts since this function is implemented on the L2/L3 border points of the network. For more information, refer to the "MAGP" page.

The two peer switches need to carry the exact same configuration of the MLAG attributes for guaranteeing proper functionality of the MLAG.

> Ensuring that both switches are configured identically is the responsibility of the user and is not monitored by the OS.

> MLAG is currently supported for 2 switches only.

> The VIP address must be on the same management IP subnet.

All nodes in an MLAG must be of the same CPU type (e.g., x86), switch type, and must all have the same OS version installed.

When working with MLAG, the maximum number of MAC addresses is limited to 88K. Without it, there is no limitation.

When transitioning from standalone into a group or vice versa, a few seconds are required for the node state to stabilize. During that time, group feature commands (e.g. MLAG commands) should not be executed. To run group features, wait for the CLI prompt to turn into [standalone:master], [<group>:master] or [<group>:standby] instead of [standalone:*unknown*] or [<group>:*unknown*].

Each MLAG VIP group must be configured with a different unicast IP address. If not, MLAG behavior is inconsistent.

In a scenario where there is no IP communication between the MGMT ports of the MLAG switches (for example when one MGMT port is disconnected), the following CLI prompt is displayed: <hostname>[<mlag cluster name>:unknown]#. This does not reflect the MLAG state, but only the state of the cluster.

IPL port-channel should not be configured with LACP rate fast, but should stay with default rate (slow).

12.9.1 MLAG Keepalive and Failover

Master election in MLAG is based on the IPs of the nodes taking part of the MLAG. The master elected is that which has the highest IPL VLAN interface local IP address.

MLAG master/slave roles take effect in fault scenarios such as split-brain, peer faults, and during software upgrades.

The MLAG pair of switches periodically exchanges a keepalive message on a user configurable interval. If the keepalive message fails to arrive for three consecutive intervals the switches break into two standalone switches. In such a case, the remaining active switch begins to act as a standalone switch and assumes that its previously peering MLAG switch has failed.

To avoid a scenario where failure on the IPL causes both MLAG peers to assume that their peer has failed, a safety mechanism is maintained based on UDP packets running via the management plane which alerts both MLAG switches that its peer is alive. In such case where keepalive packets are not received the slave shuts down its MLAG interfaces and the master becomes a standalone switch in order to avoid misalignment in MLAG configuration.

12.9.2 Unicast and Multicast Sync

Unicast and multicast sync is a mechanism which syncs the unicast and multicast FDBs of the MLAG peers. It prevents unicast asymmetric traffic from loading the network with flood traffic and multicast traffic from being processed.

12.9.3 MLAG Port Sync

Under normal circumstances, traffic from the IPL cannot pass through the MLAG ports (the IPL is isolated from the MLAG ports). If one of the MLAG links break, the other MLAG switch opens that isolation and allows traffic from its peer through the IPL to flow via the MLAG port which accesses the destination of the fallen link.

12.9.4 MLAG Virtual System-MAC

A pair of MLAG switches uses a single virtual system MAC for L2 protocols (such as LACP) operating on the MLAG ports. This virtual system MAC is served also as the STP bridge ID.

The virtual system MAC is automatically computed based on the MLAG VIP name, but can be manually set using the command "system-mac".

MLAG relies on systems to have the same virtual system MAC. Therefore, if a system MAC mismatch is detected, the slave shuts down its interfaces.

When the MLAG Virtual MAC is generated automatically (see system-mac command), the virtual MAC is calculated according to the MLAG-VIP name, using the base MAC as VRRP MAC prefix (00:00:5E:00:01:xx) with the suffix hashed from the mlag-vip name 0…255. Since there are only 256 possibilities, the Virtual MAC collisions might happen when there are many MLAG pairs in a same STP domain.

- There are only 256 possibilities when generating the virtual MAC automatically. AS such, a collisions may occur when there are many MLAG pairs in a same STP domain.
- Check for MAC collisions on every existing MLAG pair in the same STP domain for each newly added MLAG pair.

12.9.5 Upgrading MLAG Pair

Switches in the same MLAG group must have the same OS version.

When peers identify having different versions, they enter an upgrading state in which the slave peer waits for a specific period of time (according to the command "upgrade-timeout") before closing its ports.

It is advised to plan MLAG upgrade in advance and perform it in a timely manner. Do not perform topology changes during the upgrade period.

From a configuration point of view, to upgrade the MLAG cluster, first upgrade the standby switch and then upgrade the MLAG cluster master (not to be confused with the VIP master).

In the intermediate state where the standby is upgraded and the MLAG cluster master is not, traffic may be impacted by lack of synchronization between the two switches.

By default, there is 60 minutes to upgrade the MLAG cluster master. To avoid traffic impact, shorten the time between upgrading the MLAG standby and upgrading the MLAG cluster master as much as possible.

> MLAG cluster master is not to be confused with MLAG VIP master. To see MLAG VIP, run "show mlag-vip". During upgrade, the MLAG VIP master may change based on which switch is up first.
> The MLAG cluster master, on the other hand, is the switch with the highest IP address. To see which switch has a higher IP address, run "show mlag"

1. Identify the MLAG cluster master by issuing "show mlag" command. The switch with the higher local IP address of the IPL is the MLAG cluster master (in the example below, SwitchB is the master).

```
SwitchA [my-vip: master] (config)# show mlag
Admin status: Enabled
Operational status: Up
Reload-delay: 1 sec
Keepalive-interval: 30 sec
Upgrade-timeout: 60 min
System-mac: 00:00:5e:00:01:5d

MLAG Ports Configuration Summary:
Configured:  1
 Disabled:   0
 Enabled:    1

MLAG Ports Status Summary:
Inactive:        0
 Active-partial: 0
 Active-full:    1

MLAG IPLs Summary:
ID   Group         Vlan       Operational    Local        Peer        Up Time     Toggle Counter
     Port-Channel  Interface  State          IP address   IP address
-------------------------------------------------------------------------------------------------
1    Po1           1          Up             10.10.10.1      10.10.10.2     0 days      00:00:09 5
Peers state Summary:
System-id          State   Hostname
-----------------------------------
F4:52:14:2D:9B:88  Up      <SwitchA>
F4:52:14:2D:9B:08  Up       SwitchB
```

2. After identifying the MLAG master/standy, make sure to first upgrade the MLAG standby peer according to steps 1-10 in "Upgrading Operating System Software" before upgrading the master.
3. Wait until the upgraded switch is up and the "show mlag" output looks like the example is step #1, above.

   > When standby MLAG peer upgrade is complete and the master is still in the lower version, MACs are not learned by the standby switch system (except for traffic flood) until master switch upgrade is complete.

4. Upgrade the MLAG cluster master node according to steps 1-10 in "Upgrading Operating System Software".

> When two tiers of MLAG pairs are used, each pair should be upgraded sequentially and not in parallel to prevent traffic loops.

12.9.6 Interoperability with MLAG

12.9.6.1 MLAG Interoperability with L2 Protocols

MLAG inter-operates with all STP modes (RSTP, MSTP and PVRST). MLAG can be configured in a spanning tree network where the two MLAG switches function as one STP entity.

In general all static configuration must be configured identically on both peers.

Protocol	Description
Static MAC addresses	Static MAC address are not synced between MLAG peers
LACP	MPO supports all LACP modes (passive/active), but it is not a must. If used, their configuration must be identical on each peer. **Note:** if LACP system-priority is configured on one switch, and not both, it will cause MLAG port-channels to be suspended on one switch.
VLAN	VLAN membership of an MPO must be configured identically on both peers. This includes PVID, switchport mode, and tagged/untagged VLAN. VLAN static configuration such as snooping MRouter must be configured identically on both peers as well.
Spanning-tree protocol	MPO spanning-tree configuration must be identical in both switches, and other local ports' spanning-tree configuration must be done when those ports are down.
IGMP snooping	IGMP snooping must be activated globally on both peers. IGMP snooping attributes on the MPO must have identical configuration.
Port mirroring	Supported
PIM	Not supported
sFlow	Supported
LLDP	All attributes of a the MPO must be configured identically on both peers.
Isolation-groups	Not supported with MLAG
OpenFlow	Not supported over MLAG IPL
PTP	Not supported over MLAG IPL (not supported over LAG in general)
NVE	Not supported
Dot1x	Not supported

12.9.6.2 MLAG Interoperability with Routing Protocols

MLAG can operate with BGP routing protocol. For redundancy purposes, establish BGP session between the two routers over the IPL.

12.9.6.2.1 BGP over MPO

To use BGP over MPO, configure the following:

1. Configure MAGP on the relevant VLAN interface.

2. Configure IP prefix list, to match all prefixes.

```
switch (config) # ip prefix-list any seq 10 permit 0.0.0.0 /0 ge 0
```

3. Configure BGP route map to use MAGP address as nexthop. Advertised next-hop should be the virtual IP address of the MAGP.

```
switch (config) # route-map magp permit 1 match ip address any
switch (config) # route-map magp permit 1 set ip next-hop <magp-ip-address>
```

4. Apply BGP route map on neighbor egress direction.

```
switch (config) #  router bgp 65000 vrf default neighbor <neighbor> route-map magp out
```

It is not recommended to use OSPF over MLAG port-channel (MPO).

12.9.7 Configuring MLAG

This section provides a basic example of how to configure two switches and a server in an MLAG setup.

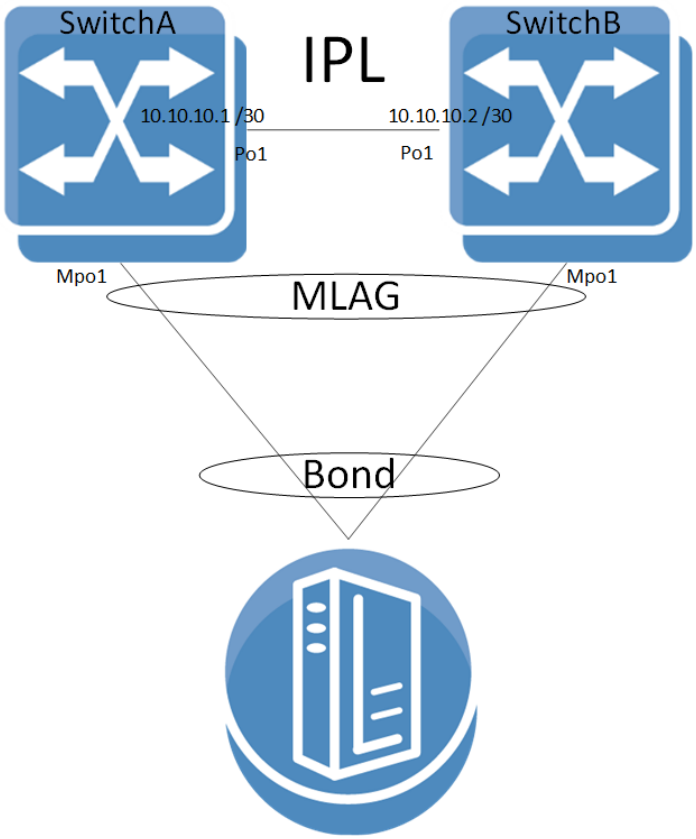

For more advanced configuration options, please refer to the community post "MLAG Procedures and Troubleshooting".

12.9.7.1 Configuring L2 MLAG

Prerequisites:

1. Enable IP routing. Run:

```
switch (config)# ip routing
```

2. (Recommended) Enable LACP in the switch. Run:

```
switch (config)# lacp
```

3. Enable the MLAG protocol commands. Run:

```
switch (config)# protocol mlag
```

Configuring the IPL:

1. Create a VLAN for the inter-peer link (IPL) to run on. Run:

```
switch (config)# vlan 4000
switch (config vlan 4000)#
```

2. Create a LAG. Run:

```
switch (config)# interface port-channel 1
switch (config interface port-channel 1)#
```

3. Map a physical port to the LAG in active mode (LACP). Run:

```
switch (config)# interface ethernet 1/1 channel-group 1 mode active
```

4. Set this LAG as an IPL. Run:

```
switch (config interface port-channel 1)# ipl 1
```

5. Create a VLAN interface. Run:

```
switch (config)# interface vlan 4000
switch (config interface vlan 4000)#
```

6. Configure MTU to 9K. Run:

```
switch (config interface vlan 4000)# mtu 9216
```

7. Set an IP address and netmask for the VLAN interface.
   Configure IP address for the IPL link on both switches:

> The IPL IP address should not be part of the management network, it could be any IP address and subnet that is not in use in the network. This address is not advertised outside the switch.

On SwitchA, run:

```
switch (config interface vlan 4000)# ip address 1.1.1.1 /30
```

On SwitchB, run:

```
switch (config interface vlan 4000)# ip address 1.1.1.2 /30
```

> The peer with the interface VLAN with the highest IP is the MLAG master. In the example, above, SwitchB (with IP 1.1.1.2) is the master. The IP addresses of both peers can be seen in via "show mlag" command.

8. Map the VLAN interface to be used on the IPL and set the peer IP address (the IP address of the IPL port on the second switch) of the IPL peer port. IPL peer ports must be configured on the same netmask.
   On SwitchA, run:

```
switch (config interface vlan 4000)# ipl 1 peer-address 1.1.1.2
```

   On SwitchB, run:

```
switch (config interface vlan 4000)# ipl 1 peer-address 1.1.1.1
```

9. (Optional) Configure a virtual IP (VIP) for the MLAG. MLAG VIP is important for retrieving peer information.

   > If you have a mgmt0 interface, the IP address should be within the subnet of the management interface. Do not use mgmt1. The management network is used for keepalive messages between the switches. The MLAG domain must be unique name for each MLAG domain. In case you have more than one pair of MLAG switches on the same network, each domain (consist of two switches) should be configured with different name.

   On SwitchA, run:

```
switch (config)# mlag-vip my-vip ip 10.234.23.254 /24
```

   On SwitchB, run:

```
switch (config)# mlag-vip my-vip
```

10. (Optional) Configure a virtual system MAC for the MLAG. Run:

```
switch (config)# mlag system-mac 00:00:5e:00:01:5d
```

Creating an MLAG interface:

1. Create an MLAG interface for the host. Run:

```
switch (config)# interface mlag-port-channel 1
switch (config interface mlag-port-channel 1)#
```

> *The MPO interfaces should be configured in the same sequence on both switches of MLAG cluster.*
>
> *Example:*
> *On SwitchA:*
> *interface mlag-port-channel 1-10*
> *interface mlag-port-channel 30-40*
> *On SwitchB:*
> *interface mlag-port-channel 1-10*
> *interface mlag-port-channel 30-40*

2. Bind an Ethernet port to the MLAG interface. Run:

```
switch (config interface ethernet 1/1)# mlag-channel-group 1 mode on
```

3. Create and enable the MLAG interface. Run:

```
switch (config interface mlag-port-channel 1)# no shutdown
```

Enabling MLAG:

1. Enable MLAG. Run:

```
switch (config mlag)# no shutdown
```

> When running MLAG as L2/L3 border point, MAGP VIP must be deployed as the default GW for MPOs. For more information, refer to "MAGP".

12.9.7.2 Verifying MLAG Configuration

1. Examine MLAG configuration and status. Run show mlag on the switch:

```
SwitchA [my-vip: master] (config)# show mlag
Admin status: Enabled
Operational status: Up
Reload-delay: 1 sec
Keepalive-interval: 30 sec
Upgrade-timeout: 60 min
System-mac: 00:00:5e:00:01:5d

MLAG Ports Configuration Summary:
Configured:  1
 Disabled:   0
 Enabled:    1

MLAG Ports Status Summary:
```

```
Inactive:        0
 Active-partial: 0
 Active-full:    1

MLAG IPLs Summary:
ID   Group       Vlan       Operational    Local        Peer        Up Time     Toggle Counter
     Port-Channel Interface  State          IP address   IP address
-------------------------------------------------------------------------------------------
1    Po1         1          Up             1.1.1.1      1.1.1.2     0 days      00:00:09 5
Peers state Summary:
System-id          State   Hostname
-----------------------------------
F4:52:14:2D:9B:88  Up      <SwitchA>
F4:52:14:2D:9B:08  Up       SwitchB
```

2. Examine the MLAG summary table. Run:

```
SwitchA [my-vip: master] (config) # show interfaces mlag-port-channel summary

MLAG Port-Channel Flags: D-Down, U-Up, P-Partial UP, S-suspended by MLAG

Port Flags:
  D: Down
  P: Up in port-channel (members)
  S: Suspend in port-channel (members)
  I: Individual

MLAG Port-Channel Summary:
  ---------------------------------------------------------------------------
  Group              Type     Local                    Peer
  Port-Channel                Ports                    Ports
  (D/U/P/S)                   (D/P/S/I)                (D/P/S/I)
  ---------------------------------------------------------------------------
  1 Mpo2(U)          Static   Eth1/2(P)                Eth1/2(P)
```

3. Examine the MLAG statistics. Run:

```
SwitchA [my-vip: master] (config)# show mlag statistics
IPL 1:
  Rx Heartbeat           : 516
  Tx Heartbeat           : 516
  Rx IGMP tunnel         : 0
  Tx IGMP tunnel         : 0
  RX XSTP tunnel         : 0
  TX XSTP tunnel         : 0
  RX mlag-notification   : 0
  TX mlag-notification   : 0
  Rx port-notification   : 0
  Tx port-notification   : 0
  Rx FDB sync            : 0
  Tx FDB sync            : 0
  RX LACP manager        : 1
  TX LACP manager        : 0
```

4. (Optional) In case MLAG-VIP was configured, its functionality can be examined using "show mlag-vip" command.

```
SwitchA [my-vip: master] (config)# show mlag-vip
MLAG VIP
========
MLAG group name: my-mlag-group
MLAG VIP address: 10.234.23.254 /24
Active nodes: 2

Hostname             VIP-State            IP Address
-----------------------------------------------------
SwitchA              master               10.234.23.1
SwitchB              standby              10.234.23.2
```

> No output will appear, if MLAG-VIP is not configured.

12.9.7.3 Enabling L3 Forwarding with User VRF

If you want to use a VRF for IP routing and forwarding on an MLAG topology, it is recommended to configure an additional VLAN interface with the same user VRF context as the non-MLAG L3 interface that has to route through the same physical ports as the IPL. This would allow forwarding L3 traffic through this VLAN interface on the same ports as the IPL.

12.9.8 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community posts:

- How To Configure MLAG
- MLAG Procedures and Troubleshooting
- Rack Solution Using SN2100 MLAG Switch Pair and ConnectX-4 Lx
- High Availability using NVIDIA Switches and Adapters
- How To Upgrade MLNX-OS Software on an MLAG Switch Pair
- How To Configure a 3rd Party Switch Connected to a Pair of MLAG Switches
- How To Enable MLAG Switch Pair Using NEO
- Configuring Cisco 6513 switch and MLAG
- How To Configure MLAG Using MLAG Wizard

12.9.9 MLAG Commands

- MLAG Commands

12.9.10 MLAG Commands

- 12.9.10.1 protocol mlag
- 12.9.10.2 mlag
- 12.9.10.3 shutdown
- 12.9.10.4 interface mlag-port-channel
- 12.9.10.5 ipl
- 12.9.10.6 ipl peer-address
- 12.9.10.7 keep-alive-interval
- 12.9.10.8 mlag-channel-group mode
- 12.9.10.9 mlag-vip
- 12.9.10.10 reload-delay
- 12.9.10.11 system-mac
- 12.9.10.12 upgrade-timeout
- 12.9.10.13 show mlag
- 12.9.10.14 show mlag-vip
- 12.9.10.15 show interfaces mlag-port-channel
- 12.9.10.16 show interfaces mlag-port-channel counters
- 12.9.10.17 show interfaces mlag-port-channel summary
- 12.9.10.18 show mlag statistics

12.9.10.1 protocol mlag

	protocol mlag no protocol mlag Enables MLAG functionality and unhides the MLAG commands. The no form of the command hides the MLAG commands and deletes its database.
Syntax Description	N/A
Default	no protocol mlag
Configuration Mode	config
History	3.3.4500
Example	`switch (config) # protocol mlag`
Related Commands	
Notes	• Running the no form of this command hides MLAG commands • MLAG may be enabled without IP routing, but without IP routing an IPL VLAN interface cannot be configured and thus MLAG does not function • MLAG may be enabled without IGMP snooping, but if IGMP snooping is disabled, multicast FDBs do not sync

12.9.10.2 mlag

	mlag Enters MLAG configuration mode.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.3.4500
Example	`switch (config) # mlag`
Related Commands	protocol mlag
Notes	

12.9.10.3 shutdown

	shutdown no shutdown Disables MLAG. The no form of the command enables MLAG.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config mlag
History	3.3.4500
Example	`switch (config mlag) # no shutdown`
Related Commands	protocol mlag

Notes	This parameter must be similar in all MLAG peers

12.9.10.4 interface mlag-port-channel

	interface mlag-port-channel <if-number> no interface mlag-port-channel <if-number> Creates an MLAG interface. The no form of the command deletes the MLAG interface.	
Syntax Description	if-number	Interface number Range: 1-1000
Default	N/A	
Configuration Mode	config	
History	3.3.4500	
Example	`switch (config) # interface mlag-port-channel 1` `switch (config interface mlag-port-channel 1) #`	
Related Commands	protocol mlag	
Notes	• The maximum number of interfaces is 64 • The default Admin state is disabled • Range configuration is possible on this interface • This interface number must be the same in all the MLAG switches	

12.9.10.5 ipl

	ipl <ipl-id> no ipl <ipl-id> Sets this LAG as an IPL port. The no form of the command resets this LAG as regular LAG.	
Syntax Description	ipl-id	IPL ID (only “1” IPL port is supported)
Default	no ipl	
Configuration Mode	config interface port-channel	
History	3.3.4500	
Example	`switch (config interface port-channel 1)# ipl 1`	
Related Commands	protocol mlag	
Notes	• If a LAG is set as IPL, only the commands “no shutdown”, “no ipl” and “no interface port-channel” become applicable • A LAG interface set as IPL must have default LAG configuration, otherwise the set is rejected. Force option can be used	

12.9.10.6 ipl peer-address

	ipl <ipl-id> peer-address <ip-address> no ipl <ipl-id> Maps a VLAN interface to be used for an IPL LAG and sets the peer IP address of the IPL peer port. The no form of the command deletes a peer IPL LAG and unbinds this VLAN interface from the IPL function.

Syntax Description	ipl-id	IPL ID (only “1” IPL port is supported)
	ip-address	IPv4 address
Default	N/A	
Configuration Mode	config interface vlan	
History	3.3.4500	
Example	`switch (config interface vlan 1)# ipl 1 peer-address 10.10.10.10`	
Related Commands	protocol mlag	
Notes	• The subnet mask is the same subnet mask of the VLAN interface • This VLAN interface should be used for IPL only	

12.9.10.7 keep-alive-interval

	keep-alive-interval <value> no keep-alive-interval Configures the interval during which keep-alive messages are issued between the MLAG switches. The no form of the command resets this parameter to its default value.	
Syntax Description	value	Time in seconds Range: 1-300
Default	1 second	
Configuration Mode	config mlag	
History	3.3.4500	
Example	`switch (config mlag) # keep-alive-interval 1`	
Related Commands	protocol mlag	
Notes	This parameter must be similar on all MLAG peers	

12.9.10.8 mlag-channel-group mode

	mlag-channel-group <if-number> mode {on | active | passive} no mlag-channel-group Binds an Ethernet port to the MLAG port-channel (MPO). The no form of the command deletes the binding.	
Syntax Description	if-number	Interface number Range: 1-1000
	on	Binds to static MLAG
	active	Sets MLAG LAG in LACP active mode
	passive	Sets MLAG LAG in LACP passive mode
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.3.4500	
Example	`switch (config interface ethernet 1/1)# mlag-channel-group 1 mode on`	
Related Commands	protocol mlag	

Notes	

12.9.10.9 mlag-vip

	mlag-vip <domain-name> ip [<ip-address> {<masklen> | netmask> [force]] no mlag-vip Sets the VIP domain and IP address for MLAG. The no form of the command deletes the VIP domain and IP address.	
Syntax Description	domain-name	MLAG group name
	<ip-address>	IP address (IPv4 only)
	<masklen>	Format example: /24 Note that a space is required between the IP address and the mask
	<netmask>	Format example: 255.255.255.0 Note that a space is required between the IP address and the mask
	force	Forces the IP address if another IP is already configured
Default	N/A	
Configuration Mode	config	
History	3.3.4500	
	3.8.2000	Updated notes
Example	`switch (config)# mlag-vip my-mlag-domain ip 10.10.10.254/24`	
Related Commands		
Notes	• This command is supported only by IPv4 address scheme. For management networks that are IPv6-only, the mlag-vip cannot be configured. • This IP address must be configured in one of the MLAG switches and must be in the box management subnet • Other switches in the MLAG must join the same domain name	

12.9.10.10 reload-delay

	reload-delay <value> no reload-delay Specifies the amount of time that MLAG ports are disabled after system reboot. The no form of the command resets this parameter to its default value.	
Syntax Description	value	Time in seconds Range: 0-300
Default	30 seconds	
Configuration Mode	config mlag	
History	3.3.4500	
Example	`switch (config mlag) # reload-delay 30`	
Related Commands		
Notes	• This interval allows the switch to learn the IPL topology to identify the master and sync the MAC address before opening the MLAG ports • This parameter must be similar in all MLAG peers	

12.9.10.11 system-mac

	system-mac <virtual-mac> no system-mac <virtual-mac> Configures virtual system MAC. The no form of the command resets this value to its default value.	
Syntax Description	virtual-mac	MAC address
Default	Default is calculated according to the MLAG-VIP name, using the base MAC as VRRP MAC prefix (00:00:5e:00:01:xx) with the suffix hashed from the mlag-vip name 0...255.	
Configuration Mode	config mlag	
History	3.4.2008	
Example	`switch (config mlag) # system-mac 00:00:5e:00:01:5d`	
Related Commands		
Notes	This parameter must be configured the same in all MLAG peers	

12.9.10.12 upgrade-timeout

	upgrade-timeout <time> no upgrade-timeout Configures the time period during which an MLAG slave keeps its ports active while in upgrading state. The no form of the command resets the parameter value to its default.	
Syntax Description	time	Time in minutes Range: 0-120 minutes
Default	60	
Configuration Mode	config mlag	
History	3.4.2008	
Example	`switch (config mlag) # upgrade-timeout 60`	
Related Commands		
Notes	This parameter must be configured the same in all MLAG peers	

12.9.10.13 show mlag

	show mlag Displays MLAG configuration and status.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4500	
	3.3.5006	Updated example
	3.4.2008	Updated example with system MAC and upgrade timeout

	3.6.6102	Updated example
Example		

```
switch (config)# show mlag

Admin status: Enabled
Operational status: Up
Reload-delay: 1 sec
Keepalive-interval: 30 sec
Upgrade-timeout: 60 min
System-mac: 00:00:5e:00:01:5d

MLAG Ports Configuration Summary:
Configured: 1
 Disabled: 0
 Enabled: 1

MLAG Ports Status Summary:

Inactive: 0
 Active-partial: 0
 Active-full: 1

MLAG IPLs Summary:
ID   Group          Vlan        Operational    Local         Peer          Up Time           Toggle
Counter
     Port-Channel   Interface   State          IP address    IP address
-------------------------------------------------------------------------------------------------
--
1    Po1            1           Up             10.10.10.1    10.10.10.2    0 days 00:00:09   5

MLAG Members Summary:
System-id          State   Hostname
-----------------------------------
f4:52:14:2d:9b:88  Up      <SwitchB>
f4:52:14:2d:9b:08  Up       SwitchA
```

Related Commands	
Notes	If run in the middle of an upgrade, the following message will appear in the output: `*Upgrading* <hostname> --> *Cluster upgrade in progress*`

12.9.10.14 show mlag-vip

	show mlag-vip Displays MLAG VIP configuration and status.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4500	
	3.6.6102	Updated example
Example		

```
switch (config) # show mlag-vip
MLAG-VIP
 MLAG group name: Test
 MLAG VIP address: 10.10.10.3/24
 Active nodes: 2

-------------------------------------------------------------
Hostname             VIP-State              IP Address
-------------------------------------------------------------
SwitchA              master                 10.10.10.1
SwitchB              standby                10.10.10.2
```

Related Commands	
Notes	

12.9.10.15 show interfaces mlag-port-channel

	show interfaces mlag-port-channel [<if-number>] Displays the MLAG LAG configuration and status.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4500	
	3.6.1002	Added “error packets” counter to Tx
	3.6.5000	Added telemetry to output
	3.6.6000	Added “forwarding mode” to output
	3.6.8008	Updated example
	3.9.1000	Updated example
Example		

```
switch (config)# show interfaces mlag-port-channel 11

Mpo11:
  Admin state         : Disabled
  Operational state   : Down
  Description         : N/A
  Mac address         : N/A
  MTU                 : 1500 bytes (Maximum packet size 1522 bytes)
  lacp-individual mode: Disabled
  Flow-control        : receive off send off
  Actual speed        : N/A
  Width reduction mode: Not supported
  Switchport mode     : access
  MAC learning mode   : Enabled
  Forwarding mode     : inherited cut-through
  FCS Ingress         : Enabled CRC check
  FCS Egress          : Enabled CRC recalculate
  FCS Timestamping    : Enabled

  Telemetry sampling: Disabled   TCs: N/A
    Telemetry threshold: Disabled        TCs: N/A
    Telemetry threshold level: N/A

  Last clearing of "show interface" counters: Never
  60 seconds ingress rate                   : 0 bits/sec, 0 bytes/sec, 0 packets/sec
  60 seconds egress rate                    : 0 bits/sec, 0 bytes/sec, 0 packets/sec

  Rx:
     0                    packets
     0                    unicast packets
     0                    multicast packets
     0                    broadcast packets
     0                    bytes
     0                    discard packets
     0                    error packets
     0                    fcs errors
     0                    undersize packets
     0                    oversize packets
     0                    pause packets
     0                    unknown control opcode
     0                    symbol errors
     0                    discard packets by storm control

  Tx:
     0                    packets
     0                    unicast packets
     0                    multicast packets
     0                    broadcast packets
     0                    bytes
     0                    discard packets
     0                    error packets
     0                    hoq discard packets
```

Related Commands	
Notes	

12.9.10.16 show interfaces mlag-port-channel counters

	show interfaces mlag-port-channel <if-number> counters Displays the extended counters for the interface.	
Syntax Description	if-numbers	MLAG interface whose properties to display
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	
	3.9.1000	Added ability to use a range of ports

Example	switch (config)# show interfaces mlag-port-channel 2-3 counters
	Mpo2:
	Rx
	  12 packets
	  0 unicast packets
	  12 multicast packets
	  0 broadcast packets
	  2700 bytes
	  0 packets of 64 bytes
	  0 packets of 65-127 bytes
	  12 packets of 128-255 bytes
	  0 packets of 256-511 bytes
	  0 packets of 512-1023 bytes
	  0 packets of 1024-1518 bytes
	  0 packets Jumbo
	  0 error packets
	  0 discard packets
	  0 fcs errors
	  0 undersize packets
	  0 oversize packets
	  0 pause packets
	  0 unknown control opcode
	  0 symbol errors
	
	Tx
	  0 packets
	  0 unicast packets
	  0 multicast packets
	  0 broadcast packets
	  152100000000 bytes
	  100000000 error packets
	  0 discard packets
	  0 pause packets
	  0 ECN marked packets
	Mpo3:
	
	. . .

Related Commands	
Notes	As of version 3.9.1000, the "if-numbers" attribute is optional. If nothing is selected, information for all ports will be displayed

12.9.10.17 show interfaces mlag-port-channel summary

	show interfaces mlag-port-channel summary Displays MLAG summary table.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4500	
	3.4.0000	Added notes and Updated example
	3.4.1100	Updated example
	3.6.6000	Updated example

Example	switch [my-vip: standby] (config)# show interfaces mlag-port-channel summary MLAG Port-Channel Flags: D-Down, U-Up, P-Partial UP, S-suspended by MLAG Port Flags: D: Down P: Up in port-channel (members) S: Suspend in port-channel (members) I: Individual MLAG Port-Channel Summary: --- Group Type Local Peer Port-Channel Ports Ports (D/U/P/S) (D/P/S/I) (D/P/S/I) --- 1 Mpo61(D) LACP Eth1/4(I) Eth1/3(S)
Related Commands	
Notes	• If a cluster is not available, the column “Peer Ports” shows “N/A”. If the cluster is available but is not configured on the peer, the “Peer Ports” column shows nothing. • If the system happens to be busy, peer ports may be unavailable and the following prompt may appear in the output: “System busy and partial information is presented - please try again later” • The “I” flag indicates an interface which is part of a LAG and in individual state • The “S” flag indicates an interface which is part of a LAG and in suspended state

12.9.10.18 show mlag statistics

	show mlag statistics Displays the MLAG IPL counters.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4500	
	3.4.0000	Updated example
	3.6.6102	Updated example
Example	switch (config)# show mlag statistics IPL 1: Rx Heartbeat : 516 Tx Heartbeat : 516 Rx IGMP tunnel : 0 Tx IGMP tunnel : 0 RX XSTP tunnel : 0 TX XSTP tunnel : 0 RX mlag-notification : 0 TX mlag-notification : 0 Rx port-notification : 0 Tx port-notification : 0 Rx FDB sync : 0 Tx FDB sync : 0 RX LACP manager : 1 TX LACP manager : 0	
Related Commands		
Notes		

12.10 Link State Tracking

A group of links may contain upstream links and downstream links. When all upstream links in a group are down, Link State Tracking (LST) shuts all the downstream links down. In order to let the peer on the other side know that it needs to stop sending traffic on the downstream links. When the upstream link recovers, LST brings up the downstream links, letting the peers know that they may resume forwarding traffic on those links.

A link can be a member of several groups. A downstream interface is shut down if at least one of the groups requests a shutdown and is brought back up if all groups request it to be up.

In situations with only downstream links in a group (no upstream links), the downstream links will stay up.

12.10.1 Configuring Link State Tracking

The following is a basic example of how to configure link state tracking group and tracking VLAN.

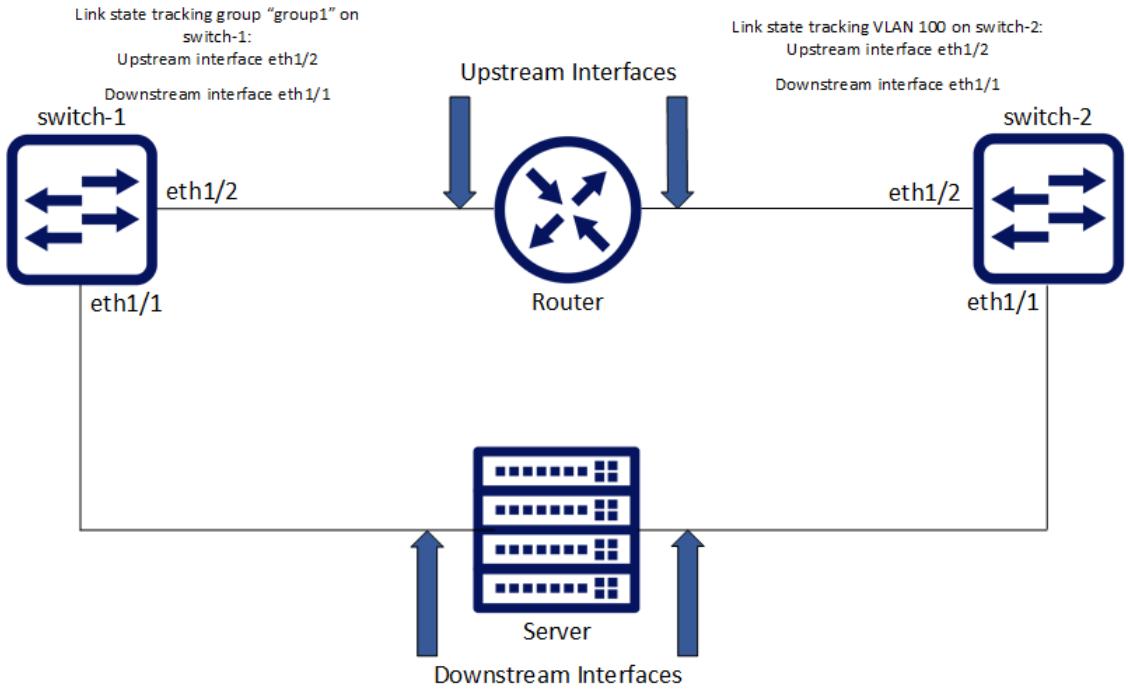

To configure Link State Tracking group:

1. Create tracking group. Run:

```
switch-1 (config) # link state tracking group group1
```

2. Configure link type on the interface. Run:

```
switch-1 (config) # interface ethernet 1/2 link type upstream
switch-1 (config) # interface ethernet 1/1 link type downstream
```

3. Add interfaces into the group. Run:

```
switch-1 (config) # interface ethernet 1/1 link state tracking group group1
switch-1 (config) # interface ethernet 1/2 link state tracking group group1
```

To configure Link State Tracking VLAN:

1. Create VLAN. Run:

```
switch-2 (config) # vlan 100
```

2. Configure VLAN members. Run:

```
switch-2 (config) # interface ethernet 1/1 switchport access vlan 100
switch-2 (config) # interface ethernet 1/2 switchport access vlan 100
```

3. Configure link type on the interface. Run:

```
switch-2 (config) # interface ethernet 1/2 link type upstream
switch-2 (config) # interface ethernet 1/1 link type downstream
```

4. Create link state tracking VLAN. Run:

```
switch-2 (config) # link state tracking vlan 100
```

To verify Link State Tracking configuration, run:

```
switch-1 (config) # show link state tracking group group1

--------------------------------------------------------------------------------
Group              Port Type      Interface          Admin Status   Operational Status
--------------------------------------------------------------------------------
group1             Upstream       Eth1/2             Enabled        Up
```

12.10.2 Link State Tracking Commands

12.10.2.1 link type

	link type {downstream | upstream} no link type Configures an interface's link direction. The no form of the command deletes the interface's link direction configuration.	
Syntax Description	downstream	Configures interface as downstream
	upstream	Configures interface as upstream
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.7.1000	
Example	`switch (config interface ethernet 1/1)# link type downstream`	
Related Commands	show link state tracking	
Notes	• IPL, loopback, and VLAN interfaces are not supported. • An interface can be either upstream or downstream but not both.	

12.10.2.2 link state tracking group

	link state tracking group <group-name> no link state tracking group <group-name> Creates a link state tracking group if one does not exist, and if applied to a specific interface, then it adds that interface to the group. The no form of the command deletes a link state tracking group, and if applied to a specific interface, then it removes that interface from the group.	
Syntax Description	group-name	Name for link state tracking group
Default	N/A	
Configuration Mode	config config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.7.1000	
Example	`switch (config interface ethernet 1/1)# link state tracking group group1`	
Related Commands	show link state tracking	
Notes	• The maximum number of tracking groups/VLANs is 64 • Link state tracking group name should not contain any of the following characters: [*/\"\\ ;,.?<>:@#$%^&()=] and should consist of no more than 255 characters • Tracking the link state of member ports in a LAG or MLAG is not supported	

12.10.2.3 link state tracking vlan

	link state tracking vlan <vlan-id> no link state tracking vlan <vlan-id> Creates a VLAN link state tracking group. All VLAN members are automatically added into this group. The no form of the command deletes a VLAN link state tracking group.	
Syntax Description	vlan-id	ID of VLAN whose link state to track
Default	N/A	
Configuration Mode	config	
History	3.7.1000	
Example	`switch (config)# link state tracking vlan 100`	
Related Commands	show link state tracking	
Notes	The maximum number of tracking groups/VLANs is 64	

12.10.2.4 show link state tracking

	show link state tracking [group <group-name> | vlan <vlan-id>] Displays link state tracking configuration.	
Syntax Description	group	Displays link state tracking per tracking group
	vlan	Displays link state tracking per VLAN
Default	N/A	
Configuration Mode	Any command mode	

History	3.7.1000
Example	

```
switch (config)# show link state tracking

------------------------------------------------------------------------------------
Group               Port Type      Interface          Admin Status   Operational Status
------------------------------------------------------------------------------------
Vlan 100            Upstream       Eth1/54            Enabled        Down
Vlan 100            Downstream     Eth1/1             Enabled        Down (by tracking)
Vlan 100            Unassigned     Eth1/2             Enabled        Up
Vlan 101            Upstream       Eth1/54            Enabled        Down
Vlan 101            Downstream     Eth1/1             Enabled        Down (by tracking)
Vlan 101            Unassigned     Eth1/2             Enabled        Up
group1              Downstream     Eth1/1             Enabled        Down (by tracking)
```

Related Commands	link type link state tracking group link state tracking vlan
Notes	The maximum number of tracking groups/VLANs is 64

12.11 QinQ

A QinQ VLAN tunnel enables a service provider (SP) to segregate the traffic of different customers in their infrastructure, while still giving the customer a full range of VLANs for their internal use by adding a second 802.1Q VLAN tag to an already tagged frame.

So let us assume for example that an SP exists which needs to offer L2 connectivity to two corporations, “X” and “Y”, that have campuses located in both “A”, “B”. All campuses run Ethernet LANs, and the customers intend to connect through the SP’s L2 VPN network so that their campuses are in the same LAN (L2 network). Hence, it would be desirable for “X”, “Y” to have a single LAN each in both “A”, “B” which could easily exceed the VLAN limit of 4096 of the 802.1Q specification.

12.11.1 QinQ Operation Modes

QinQ can be enabled on a port or according to predefined conditions.

C-VLAN is the VLAN tag assigned to the ingress traffic of a QinQ-enabled interface.

S-VLAN is the VLAN tag assigned to the egress traffic of a QinQ-enabled interface.

- ACL-mode: Adding and removing S-VLAN is determined by an ACL-dependent action
- Port-mode: All ingress traffic to a specific QinQ-enabled interface is tagged with an additional VLAN 802.1Q tag (also known as S-VLAN). The S-VLAN ID is equal to that interface’s PVID (access VLAN).
  The S-VLAN tag is added regardless of whether the traffic is tagged or untagged. Traffic coming out from this port, has the S-VLAN stripped from it.

12.11.2 Configuring QinQ

1. Create the C-VLAN. Run:

```
switch (config) # vlan 200
switch (config vlan 200) # exit
```

2. Enter the configuration mode of an Ethernet, LAG, or MLAG interface. Run:

```
switch (config) # interface port-channel 100
```

3. Change the switchport mode of the interface to enable QinQ. Run:

```
switch (config interface port-channel 100) # switchport mode dot1q-tunnel
```

4. Change its port VLAN ID (PVID). This configures the S-VLAN. Run:

```
switch (config interface port-channel 100) # switchport access vlan 200
```

5. Verify the configuration. Run:

```
switch (config interface port-channel 100) # show interface port-channel 100

Po100
  Admin state: Enabled
  Operational state: Up
  Description: N\A
  Mac address: 00:00:00:00:00:00
    MTU: 1500 bytes(Maximum packet size 1522 bytes)
  lacp-individual mode: Disabled
  Flow-control: receive off send off
  Actual speed: 1 X 40 Gbps
  Width reduction mode: disabled
  Switchport mode: dot1q-tunnel
  QoS mode: uniform
  MAC learning mode: Enabled
  Last clearing of "show interface" counters : Never
  60 seconds ingress rate: 0 bits/sec, 0 bytes/sec, 0 packets/sec
  60 seconds egress rate: 0 bits/sec, 0 bytes/sec, 0 packets/sec

Rx
  0 packets
  0 unicast packets
  0 multicast packets
  0 broadcast packets
  0 bytes
  0 error packets
  0 discard packets

Tx
  0 packets
  0 unicast packets
  0 multicast packets
  0 broadcast packets
  0 bytes
  0 discard packets
```

6. Verify the configuration. Run:

```
switch (config interface port-channel 100) # show interfaces switchport
Interface       Mode          Access vlan       Allowed vlans
-------------------------------------------------------------------------
Eth1/1          access        1
Eth1/2          access        1
Eth1/3          access        1
Eth1/4          access        1
Eth1/5          access        1
Eth1/6          access        1
...
Eth1/27         access        1
Eth1/33         access        1
Eth1/34         access        1
Eth1/35         access        1
Eth1/36         access        1
Po400           dot1q-tunnel  200
```

12.11.3 QinQ Commands

12.11.3.1 switchport dot1q-tunnel qos-mode

	switchport dot1q-tunnel qos-mode {pipe | uniform} no switchport dot1q-tunnel qos-mode Assigns QoS to the service provider's traffic. The no form of the command resets the parameter value to its default.	
Syntax Description	pipe	Gives the service provider's traffic the same QoS as the customer's traffic
	uniform	Gives the service provider's traffic QoS 0
Default	pipe	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.4.3000	
Example	`switch (config interface ethernet 1/1) # switchport dot1q-tunnel qos-mode uniform`	
Related Commands	show vlan show interfaces switchport switchport access vlan switchport [trunk | hybrid] allowed-vlan vlan	
Notes		

12.12 Access Control List (ACL)

An Access Control List (ACL) is a list of permissions attached to an object, to filter or match switches packets. When the pattern is matched at the hardware lookup engine, a specified action (e.g. permit/deny) is applied. The rule fields represent flow characteristics such as source and destination addresses, protocol and VLAN ID.

ACL support currently allows actions of permit or deny rules, with or without mirroring, and supports only ingress direction. ACL search pattern can be taken from either L2 or L3 fields, e.g L2/ L3 source and destination addresses, protocol, VLAN ID and priority or TCP port.

12.12.1 Configuring ACL

ACL is configured by the user and is applied to a port once the ACL search engine matches search criteria with a received packet.

To configure ACL:

1. Create a MAC / IPv4 ACL (access-list) entity. Run:

```
switch (config) mac access-list mac-acl
switch (config mac access-list mac-acl) #
```

2. Add a MAC / IP rules to the appropriate access-list. Run:

```
switch (config mac access-list mac-acl) # seq-number 10 deny 0a:0a:0a:0a:0a:0a mask ff:ff:ff:ff:ff:ff any
vlan 6 cos 2 protocol 80
```

3. Bind the created access-list to an interface (port or LAG). Run:

```
switch (config) # interface ethernet 1/1
switch (config interface ethernet 1/1) # mac port access-group mac-acl
```

12.12.2 ACL Actions

An ACL action is a set of actions can be activated in case the packet hits the ACL rule.

To modify the VLAN tag of the egress traffic as part of the ACL “permit” rule:

1. Create access-list action profile:
    a. Create an action access-list profile using the command “access-list action <action-profile-name>”.
    b. Add rule to map a VLAN using the command “vlan-map <vlan-id>” within the action profile configuration mode.
    c. Add action on a rule to strip the VLAN from a packet using the command “vlan-pop” within the action profile configuration mode.
    d. Add action on a rule to append a VLAN to a packet using the command “vlan-push” within the action profile configuration mode.
2. Create an access-list and bind the action rule:
    a. Create an access-list profile using the command “{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list”.
    b. Add access list rule using the command “deny/permit” (“action <action profile name>”).
3. Bind the access-list to an interface using the command “{ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group”.

```
Create an action profile and add vlan mapping action:
switch (config)# access-list action my-action
switch (config access-list action my-action)# vlan-map 20
switch (config access-list action my-action)# exit

Create an access list and bind rules:
switch (config)# mac access-list my-list
switch (config mac access-list my-list)# permit any any action my-action
switch (config mac access-list my-list)# exit

Bind an access-list to a port:
switch (config)# interface ethernet 1/1
switch (config interface ethernet 1/1)# mac access-list my-list
```

To mirror traffic to the monitor session as part of the ACL “permit” rule"

1. Create access-list action profile:
    a. Create an action access-list profile using the command “access-list action ”.
    b. Add a rule to mirror traffic to monitor session using the command “monitor session” within the action profile configuration mode.
2. Create an access-list and bind the action rule:
    a. Create an access-list profile using the command “{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list”.

b. Add access list rule using the command “deny/permit” (“action ”).

3. Bind the access-list to an interface using the command “{ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group

```
Create an action profile and add monitor mapping action:
switch (config)# access-list action my-action
switch (config access-list action my-action)# monitor session 1
switch (config access-list action my-action)# exit

Create an access list and bind rules:
switch (config)# mac access-list my-list
switch (config mac access-list my-list)# permit any any vlan 10 action my-action
switch (config mac access-list my-list)# exit

Bind an access-list to a port:
switch (config)# interface ethernet 1/1
switch (config interface ethernet 1/1)# mac access-list my-list
switch (config interface ethernet 1/1)# exit
```

12.12.3 ACL Logging

A strong insight into the system is given by ACL logging. ACLs can log packets that pass through the switch, so the flows can later be analyzed.

A packet that hits an ACL with a log clause is passed to the logger. The logger writes the partial header of the packet (L2 or L3) to the syslog, with a timestamp and some additional information such as ingress interface and the VLAN to which the packet belongs.

To protect the system memory, a limited number of flows are collected for each time interval. If the number of flows for a specific time interval is exceeded, then no packets are logged for this time interval.

To further protect the system, a rate-limiter controls the number of packets passed to the CPU.

> Only packets traversing the switch are logged. Packets that are passed to the CPU are not.

12.12.4 ACL Capability Summary

The following table summarizes the ACL capabilities supported by NVIDIA Onyx.

ACL Table	Policy	Protocol	Keys	Actions	Supported Interfaces (Ingress Bind Point Only)
MAC	Permit Deny Remark	N/A	DST MAC (with mask) SRC MAC (with mask) Protocol CoS VLAN-ID VLAN-group	VLAN map VLAN pop VLAN push Counter per rule Shared counter to rules Log Policer Mirroring	L2 port LAG MLAG RIF VLAN interface

ACL Table	Policy	Protocol	Keys	Actions	Supported Interfaces (Ingress Bind Point Only)
IPv4	Permit Deny Remark	IP	DST IP (incl. subnets) SRC IP (incl. subnets)	VLAN map VLAN pop VLAN push Counter per rule Shared counter to rules Log Policer Mirroring	L2 port LAG MLAG RIF VLAN interface
		TCP	DST IP (incl. subnets) SRC IP (incl. subnets) L4 DST port (incl. range) L4 SRC port (incl. range) TCP flags Establish flow		
		UDP	DST IP (incl. subnets) SRC IP (incl. subnets) L4 DST port (incl. range) L4 SRC port (incl. range)		
		TCP-UDP	DST IP (incl. subnets) SRC IP (incl. subnets) L4 DST port (incl. range) L4 SRC port (incl. range)		
		ICMP	DST IP (incl. subnets) SRC IP (incl. subnets) Code Type		
IPv6	Permit Deny Remark	IPv6	DST IPv6 (incl. subnets) SRC IPv6 (incl. subnets)	VLAN map VLAN pop VLAN push Counter per rule Shared counter to rules Log Policer Mirroring	L2 port LAG MLAG RIF VLAN interface
		TCP	DST IPv6 (incl. subnets) SRC IPv6 (incl. subnets) L4 DST port (incl. range) L4 SRC port (incl. range) TCP flags Establish flow		
		UDP	DST IPv6 (incl. subnets) SRC IPv6 (incl. subnets) L4 DST port (incl. range) L4 SRC port (incl. range)		

ACL Table	Policy	Protocol	Keys	Actions	Supported Interfaces (Ingress Bind Point Only)
		TCP-UDP	DST IPv6 (incl. subnets) SRC IPv6 (incl. subnets) L4 DST port (incl. range) L4 SRC port (incl. range)		
		ICMPv6	DST IPv6 (incl. subnets) SRC IPv6 (incl. subnets) Code Type		
MAC-UDK	Permit Deny Remark	N/A	DST MAC (with mask) SRC MAC (with mask) Protocol CoS VLAN-ID VLAN-group UDK1 (up to 4 bytes) UDK2 (up to 4 bytes) UDK3 (up to 4 bytes) UDK4 (up to 4 bytes)	VLAN map VLAN pop VLAN push Counter per rule Shared counter to rules Log Policer Mirroring	L2 port LAG MLAG RIF VLAN interface
IPv4-UDK	Permit Deny Remark	IP	DST IP (incl. subnets) SRC IP (incl. subnets) UDK1 (up to 4 bytes) UDK2 (up to 4 bytes) UDK3 (up to 4 bytes) UDK4 (up to 4 bytes)	VLAN map VLAN pop VLAN push Counter per rule Shared counter to rules Log Policer Mirroring	L2 port LAG MLAG RIF VLAN interface
		TCP	DST IP (incl. subnets) SRC IP (incl. subnets) L4 DST port (incl. range) L4 SRC port (incl. range) TCP flags Establish flow UDK1 (up to 4 bytes) UDK2 (up to 4 bytes) UDK3 (up to 4 bytes) UDK4 (up to 4 bytes)		
		UDP	DST IP (incl. subnets) SRC IP (incl. subnets) L4 DST port (incl. range) L4 SRC port (incl. range) UDK1 (up to 4 bytes) UDK2 (up to 4 bytes) UDK3 (up to 4 bytes) UDK4 (up to 4 bytes)		

ACL Table	Policy	Protocol	Keys	Actions	Supported Interfaces (Ingress Bind Point Only)
		TCP-UDP	DST IP (incl. subnets) SRC IP (incl. subnets) L4 DST port (incl. range) L4 SRC port (incl. range) UDK1 (up to 4 bytes) UDK2 (up to 4 bytes) UDK3 (up to 4 bytes) UDK4 (up to 4 bytes)		
		ICMP	DST IP (incl. subnets) SRC IP (incl. subnets) Code Type UDK1 (up to 4 bytes) UDK2 (up to 4 bytes) UDK3 (up to 4 bytes) UDK4 (up to 4 bytes)		

*The maximum number of rules that can be configured per ACL type depends on the system resources utilized by the existing configuration. In order to reach the maximum number of rules, as defined in the table above, disable IP routing.

12.12.5 Additional Readings and Use Cases

For more information about this feature and its potential applications, please refer to the following community post:

- HowTo Configure Filtering Rules on Ethernet Switches (ACLs, IP Filtering)

12.12.6 ACL Commands

- ACL Commands

12.12.7 ACL Commands

- 12.12.7.1 {ipv4/ipv6/mac/ipv4-udk/mac-udk} access-list
- 12.12.7.2 policer
- 12.12.7.3 bind-point rif
- 12.12.7.4 remark
- 12.12.7.5 shared-counter
- 12.12.7.6 clear shared-counters
- 12.12.7.7 clear counters

- 12.12.7.8 {ipv4/ipv6/mac/ipv4-udk/mac-udk} access-list clear counters
- 12.12.7.9 {ipv4/ipv6/mac/ipv4-udk/mac-udk} port access-group
- 12.12.7.10 deny/permit (MAC ACL rule)
- 12.12.7.11 deny/permit (IPv4 ACL rule)
- 12.12.7.12 deny/permit (IPv4 TCP ACL rule)
- 12.12.7.13 deny/permit (IPv4 TCP-UDP/UDP ACL rule)
- 12.12.7.14 deny/permit (IPv4 ICMP ACL rule)
- 12.12.7.15 deny/permit (IPv6 ACL rule)
- 12.12.7.16 deny/permit (IPv6 TCP ACL rule)
- 12.12.7.17 deny/permit (IPv6 TCP-UDP/UDP ACL rule)
- 12.12.7.18 deny/permit (IPv6 ICMPv6 ACL rule)
- 12.12.7.19 deny/permit (MAC UDK ACL rule)
- 12.12.7.20 deny/permit (IPv4 UDK ACL rule)
- 12.12.7.21 deny/permit (IPv4 TCP UDK ACL rule)
- 12.12.7.22 deny/permit (IPv4 TCP-UDP/UDP UDK ACL rule)
- 12.12.7.23 deny/permit (IPv4 ICMP UDK ACL rule)
- 12.12.7.24 port access-group (IPv4/IPv4 UDK/IPv6/MAC/MAC UDK)
- 12.12.7.25 access-list action
- 12.12.7.26 access-list log
- 12.12.7.27 vlan-map
- 12.12.7.28 vlan-pop
- 12.12.7.29 vlan-push
- 12.12.7.30 monitor session
- 12.12.7.31 show ipv4 access-lists
- 12.12.7.32 show ipv4-udk access-lists
- 12.12.7.33 show ipv6 access-lists
- 12.12.7.34 show mac access-lists
- 12.12.7.35 show mac access-lists summary
- 12.12.7.36 show mac-udk access-lists
- 12.12.7.37 show access-lists action
- 12.12.7.38 show mac-udk access-lists
- 12.12.7.39 show access-lists log config
- 12.12.7.40 show access-lists policers (ipv4/ipv4-udk/ipv6/mac/mac-udk)
- 12.12.7.41 show access-lists shared-counters (ipv4/ipv4-udk/ipv6/mac/mac-udk)
- 12.12.7.42 show access-lists summary
- 12.12.7.43 show access-lists log
- 12.12.7.44 show access-lists log config

12.12.7.1 {ipv4/ipv6/mac/ipv4-udk/mac-udk} access-list

	{ipv4 | ipv6 | mac | ipv4-udk | mac-udk} access-list <acl-name> no {ipv4 | ipv6 | mac | ipv4-udk | mac-udk} access-list <acl-name> Creates an ACL table and enters its configuration mode. The no form of the command deletes the ACL table.	
Syntax Description	ipv4 | mac	IPv4 or MAC – access list
	acl-name	User-defined string for the ACL

Default	No ACL available by default.	
Configuration Mode	config	
History	3.1.1400	
	3.6.5000	Added ipv6, ipv4-udk, and mac-udk parameters
Example	`switch (config)# mac access-list my-mac-list` `switch (config mac access-list my-mac-list)#`	
Related Commands	ipv4/port access-group	
Notes	• Each table has its own set of predefined keys • The mac-udk and ipv4-udk options add an extra UDK to the standard MAC and IPv4 tables • When a new access-list is created, its default bind port is L2 port	

12.12.7.2 policer

	policer <policer_name> {bits|bytes|packets} rate <rate_value> [k|m|g] [burst <burst_value> [k|m|g]] no policer <policer_name> Creates a new shared-policer that can be bound to rules on this table. The no form of the command removes the policer	
Syntax Description	rate_value	Policer rate value (of the bits, bytes, or packets) Default is bits
	burst_value	Sets burst to policer. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000. For bits there is no default burst. Min value: 2000 bytes.
	k, m, g	Rate/burst value units: kilo, mega, or giga—not mandatory.
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer. Min value: 8000 bits.
	packets	Attaches packets type policer
	rate	Policer rate value: 100-1000000000000
Default	Disabled	
Configuration Mode	config mac access-list config ipv4 access-list config ipv6 access-list config ipv4-udk access-list config mac-udk access-list	
History	3.6.5000	
Example	`switch (config mac access-list my-mac-list) # policer myPolicer packets rate 1000`	
Related Commands	ipv4/ipv6/mac/ipv4-udk/mac-udk access-list	
Notes	• This ACL policer is shared when this table is bound to two or more ports. • The policer configuration will always be displayed in bytes	

12.12.7.3 bind-point rif

	bind-point rif no bind-point rif Changes the ACL table bind point from L2 port mode to L3 port. The no form of the command resets this parameter to its default.
Syntax Description	N/A
Default	L2 port
Configuration Mode	config mac access-list config ipv4 access-list config ipv6 access-list config ipv4-udk access-list config mac-udk access-list
History	3.6.5000
Example	`switch (config mac access-list my-mac-list)# bind-point rif`
Related Commands	ipv4/ipv6/mac/ipv4-udk/mac-udk access-list
Notes	• The bind point may only be changed when an ACL table is empty (no rules) and unbound • This command is used to attach ACLs to interface VLANs only

12.12.7.4 remark

	[<seq-number>] remark <string> no [<seq-number>] remark <string> Creates a remark rule from an ACL table. The no form of the command deletes a remark rule from an ACL table.
Syntax Description	N/A
Default	N/A
Configuration Mode	config mac access-list config ipv4 access-list config ipv6 access-list config ipv4-udk access-list config mac-udk access-list
History	3.6.5000
Example	`switch (config mac access-list my-mac-list)# remark "1st group"`
Related Commands	ipv4/ipv6/mac/ipv4-udk/mac-udk access-list
Notes	• The remark rule has a sequence number like standard rules and it can be displayed when showing all rules of ACL table • This rule has no effect on traffic and it is only for management purposes

12.12.7.5 shared-counter

	shared-counter <counter-name> no shared-counter <counter-name> Creates a shared counter. The no form of the command deletes a shared counter.

Syntax Description	counter-name	Shared counter name
Default	N/A	
Configuration Mode	config mac access-list config ipv4 access-list config ipv6 access-list config ipv4-udk access-list config mac-udk access-list	
History	3.6.5000	
Example	`switch (config mac access-list my-mac-list)# shared-counter myCounter`	
Related Commands	ipv4/ipv6/mac/ipv4-udk/mac-udk access-list	
Notes	• When creating a new shared counter, it is created only in the scope of the ACL table it has been initially created on and cannot be shared across multiple ACL tables • A shared counter cannot be deleted when attached to rules	

12.12.7.6 clear shared-counters

	clear shared-counters [<counter-name>] Resets all shared counters in ACL table or a specific shared counter.	
Syntax Description	counter-name	Shared counter name
Default	N/A	
Configuration Mode	config mac access-list config ipv4 access-list config ipv6 access-list config ipv4-udk access-list config mac-udk access-list	
History	3.6.5000	
Example	`switch (config mac access-list my-mac-list)# clear shared-counters`	
Related Commands	ipv4/ipv6/mac/ipv4-udk/mac-udk access-list shared-counter	
Notes		

12.12.7.7 clear counters

	clear counters [<seq-number>] Resets all counters (including shared counters) in ACL table or a specific counter.	
Syntax Description	seq-number	The sequence number of the rule whose counter to reset
Default	N/A	
Configuration Mode	config mac access-list config ipv4 access-list config ipv6 access-list config ipv4-udk access-list config mac-udk access-list	
History	3.6.5000	
Example	`switch (config mac access-list my-mac-list)# clear counters 10`	

Related Commands	ipv4/ipv6/mac/ipv4-udk/mac-udk access-list shared-counter
Notes	

12.12.7.8 {ipv4/ipv6/mac/ipv4-udk/mac-udk} access-list clear counters

	{ipv4 | ipv6 | mac | ipv4-udk | mac-udk} access-list clear counters Resets all counters (including shared counters) on all ACL tables of the same type.
Syntax Description	N/A
Default	N/A
Configuration Mode	config mac access-list config ipv4 access-list config ipv6 access-list config ipv4-udk access-list config mac-udk access-list
History	3.6.5000
Example	`switch (config)# ipv4 access-list clear counters`
Related Commands	ipv4/ipv6/mac/ipv4-udk/mac-udk access-list shared-counter
Notes	

12.12.7.9 {ipv4/ipv6/mac/ipv4-udk/mac-udk} port access-group

	{ipv4 | ipv6 | mac | ipv4-udk | mac-udk} port access-group <acl-name> no {ipv4 | ipv6 | mac | ipv4-udk | mac-udk} port access-group <acl-name> Binds an ACL to the interface. The no form of the command unbinds the ACL from the interface.	
Syntax Description	ipv4 | mac	IPv4 or MAC - access list
	acl-name	ACL name
Default	No ACL is bind by default.	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel config interface vlan	
History	3.1.1400	
	3.3.4500	Added MPO configuration mode
	3.6.5000	Added new parameters
Example	`switch (config interface ethernet 1/1) # mac port access-group my-list`	
Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list	
Notes	The access control list should be defined prior to the binding action	

12.12.7.10 deny/permit (MAC ACL rule)

	[seq-number <sequence-number>] {permit | deny} ip {<source-mac> mask <mac_mask> | any} {<dest-mac> mask <mac_mask> | any} [protocol <protocol_num>] [cos <cos>] [vlan <vlan_id>] [vlan-mask <vlan_mask>] [action <action-name>] [log] [counter | shared-counter <name>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]} no <sequence-number> Creates a rule for MAC ACL. The no form of the command deletes a rule from the MAC ACL.	
Syntax Description	sequence-number	Optional parameter to set a specific sequence number for the rule Range: 1-65535
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-mac> mask <mac_mask> | any	Sets source MAC and optionally sets a mask for that MAC. The “any” option will cause the rule not to check the source MAC.
	<dest-mac> mask <mac_mask> | any	Sets destination MAC and optionally sets a mask for that MAC. The “any” option will cause the rule not to check the destination MAC.
	protocol	Sets the Ethertype field value from the MAC address Range: 0x0000-0xffff
	cos	Sets the COS (priority bit) field Range: 0-7
	vlan <vlan_id>	Sets the VLAN ID field Range: 1-4094
	vlan-mask <vlan-mask>	Sets VLAN group Range: 0x0000-0x0FFF
	action	Action name (free string)
	log	Enable the log option
	counter	Attach a unique counter to rule
	shared-counter	Attach a predefined shared-counter to rule
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer. Min value: 8000 bits.
	packets	Attaches packets type policer
	rate	Policer rate value: 100-1000000000000
	k | m | g	Specifies kilo, mega, giga
	burst	Sets burst to policer. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000. For bits there is no default burst. Min value: 2000 bytes.

	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority Range: 0-7
	tc <tc_value>	Mapping of matched traffic to TC Range: 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config mac acl	
History	3.1.1400	
	3.3.4500	Added vlan-mask parameter
	3.5.1000	Updated seq-number parameter
	3.6.5000	Added log, counter, and shared-counter parameters
	3.6.6000	Added policer parameters
	3.7.0000	Added bits, switch-priority and tc parameters
Example	`switch (config mac access-list my-list) # seq-number 10 deny 0a:0a:0a:0a:0a:0a mask ff:ff:ff:ff:ff:ff any vlan 6 cos 2 protocol 80`	
Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group	
Notes	• VLAN and VLAN group cannot be used in the same command • It is possible to attach the rule to a unique policer, or to create a policer only for the rule • The policer configuration will always be displayed in bytes • This ACL policer is shared when this table is bound to two or more ports.	

12.12.7.11 deny/permit (IPv4 ACL rule)

	[seq-number <sequence-number>] {permit | deny} ip {<source-ip> mask <ip> | [any]} {<dest-ip> mask <ip> | [any]} [action <action-id>] [log] [counter | shared-counter <name>] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]} no <sequence-number> Creates a rule for IPv4 ACL. The no form of the command deletes a rule from the IPv4 ACL.	
Syntax Description	sequence-number	Optional parameter to set a specific sequence number for the rule Range: 1-65535
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	{any | <source-ip> mask <ip>}	Sets source IP and optionally sets a mask for that IP address. The “any” option causes the rule to not check the source IP. Range: 0-255.
	{any | <destination-ip> mask <ip>}	Sets destination IP and optionally sets a mask for that IP. The “any” option causes the rule to not check the destination IP.

	action	Action needs to be defined before attaching to rule
	log	Enable the log option
	counter	Attach a unique counter to rule
	shared-counter	Attach a predefined shared-counter to rule
	ecn	ECN ACL filter Range: 0-3
	ttl	Time to live ACL filter Range: 0-3
	dscp	DSCP ACL filter Range: 0-63
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer. Min value: 8000 bits.
	packets	Attaches packets type policer
	rate	Policer rate value: 100-1000000000000
	k | m | g	Specifies kilo, mega, giga
	burst	Sets burst to policer. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000. For bits there is no default burst. Min value: 2000 bytes.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority Range: 0-7
	tc <tc_value>	Mapping of matched traffic to TC Range: 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv4 acl	
History	3.1.1400	
	3.3.4302	Updated syntax description of mask <ip> parameter
	3.5.1000	Updated seq-number parameter
	3.6.5000	Added log, counter, and shared-counter parameters
	3.6.6000	Added ECN, TTL, DSCP, and policer parameters
	3.7.0000	Added bits, switch-priority, and tc parameters
Example	`switch (config ipv4 access-list my-list) # deny ip any any action act shared-counter`	
Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group	

Notes	• User cannot attach a shared counter defined on a different ACL table • The parameter shared-counter must be defined before attaching it to the scope of the ACL table • It is possible to attach the rule to a unique policer, or to create a policer only for the rule • The policer configuration will always be displayed in bytes • This ACL policer is shared when this table is bound to two or more ports.

12.12.7.12 deny/permit (IPv4 TCP ACL rule)

	[seq-number <sequence-number>] {deny | permit} tcp {<source-ip> mask <ip> | any} {<dest-ip> mask <ip> | any} [src-port <src-port> | eq-source <src-port> | src-port-range <from> <to>] [dest-port <dest-port> | eq-destination <dest-port> | dest-port-range <from> <to>] [action <action-id>] [established | [ack {0 | 1}] [urg {0 | 1}] [rst {0 | 1}] [syn {0 | 1}] [fin {0 | 1}] [psh {0 | 1}] [ns {0 | 1}] [ece {0 | 1}] [cwr {0 | 1}]] [log] [counter | shared-counter <name>] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]} no <sequence-number> Creates a rule for IPv4 TCP ACL. The no form of the command deletes a rule from the ACL.	
Syntax Description	sequence-number	Optional parameter to set a specific sequence number for the rule Range: 1-65535
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-ip> mask <ip> | any	Sets source IP and optionally sets a mask for that IP address. The “any” option will cause the rule not to check the source IP.
	<dest-ip> mask <ip> | any	Sets destination IP and optionally sets a mask for that IP. The “any” option will cause the rule not to check the destination IP.
	src-port	L4 source port Note: User may only choose one of the following options to configure source port: src-port; eq-source
	eq-source <src-port>	TCP source port number Range: 0-65535
	src-port-range	Sets a range of L4 source ports to match Note: User may configure either a single source port or a range
	dest-port	L4 destination port Note: User may only choose one of the following options to configure destination port: dest-port; eq-destination
	eq-destination <dest-port>	TCP destination port number Range: 0-65535
	dest-port-range	Sets a range of L4 destination ports to match Note: User may configure either a single destination port or a range
	action	Action needs to be defined before attaching to rule
	established	Matches flows which are in established state (“ack” or “rst” flags are set)

	ack; urg; rst; syn; fin; psh; ns; ece; cwr	Matches flows with specific flag Possible match: 0 or 1
	log	Enables the log option
	counter	Attaches a unique counter to rule
	shared-counter	Attaches a predefined shared-counter to rule
	ecn	ECN ACL filter Range: 0-3
	ttl	Time to live ACL filter Range: 0-225
	dscp	DSCP ACL filter Range: 0-63
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer. Min value: 8000 bits.
	packets	Attaches packets type policer
	rate	Policer rate value Range: 100-1000000000000
	k | m | g	Specifies kilo, mega, giga
	burst	Sets burst to policer. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000. For bits there is no default burst. Min value: 2000 bytes.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority Range: 0-7
	tc <tc_value>	Mapping of matched traffic to TC Range: 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv4 acl	
History	3.1.1400	
	3.5.1000	Updated seq-number parameter
	3.6.5000	Updated command syntax
	3.6.6000	Added ECN, TTL, DSCP, policer, and extra flag parameters
	3.7.0000	Added bits, switch-priority and tc parameters
Example	`switch (config ipv4 access-list my-list)# permit tcp any any src-port 200 dest-port-range 200 400 established` `switch (config ipv4 access-list my-list)# permit tcp any any ns 0 policer packets rate 1 k burst 2050`	

Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group
Notes	• L4 ports are valid • It is possible to attach the rule to a unique policer, or to create a policer only for the rule • The policer configuration will always be displayed in bytes • This ACL policer is shared when this table is bound to two or more ports.

12.12.7.13 deny/permit (IPv4 TCP-UDP/UDP ACL rule)

	[seq-number <sequence-number>] {deny | permit} {tcp-udp | udp} {<source-ip> mask <ip> | any} {<dest-ip> mask <ip> | any} [src-port <src-port> | eq-source <src-port> | src-port-range <from> <to>] [dest-port <dest-port> | eq-destination <dest-port> | dest-port-range <from> <to>] [action <action-id>] [log] [counter | shared-counter <name>] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]} no <sequence-number> Creates a rule for IPv4 TCP-UDP/UDP ACL. The no form of the command deletes a rule from the ACL.	
Syntax Description	sequence-number	Optional parameter to set a specific sequence number for the rule Range: 1-65535
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-ip> mask <ip> | any	Sets source IP and optionally sets a mask for that IP address. The “any” option will cause the rule not to check the source IP.
	<dest-ip> mask <ip> | any	Sets destination IP and optionally sets a mask for that IP. The “any” option will cause the rule not to check the destination IP.
	src-port	L4 source port Note: User may only choose one of the following options to configure source port: src-port; eq-source
	eq-source <src-port>	TCP-UDP/UDP source port number Range: 0-65535
	src-port-range	Sets a range of L4 source ports to match Note: User may configure either a single source port or a range
	dest-port	L4 destination port Note: User may only choose one of the following options to configure destination port: dest-port; eq-destination
	eq-destination <dest-port>	TCP-UDP/UDP destination port number Range: 0-65535
	dest-port-range	Sets a range of L4 destination ports to match Note: User may configure either a single destination port or a range
	action	Action needs to be defined before attaching to rule
	log	Enables the log option
	counter	Attaches a unique counter to rule
	shared-counter	Attaches a predefined shared-counter to rule

	ecn	ECN ACL filter Range: 0-3
	ttl	Time to live ACL filter Range: 0-225
	dscp	DSCP ACL filter Range: 0-63
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer. Min value: 8000 bits.
	packets	Attaches packets type policer
	rate	Policer rate value Range: 100-1000000000000
	k | m | g	Specifies kilo, mega, giga
	burst	Sets burst to policer. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000. For bits there is no default burst. Min value: 2000 bytes.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority Range: 0-7
	tc <tc_value>	Mapping of matched traffic to TC Range: 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv4 acl	
History	3.1.1400	
	3.5.1000	Updated seq-number parameter
	3.6.5000	Updated command syntax
	3.6.6000	Added ECN, TTL, DSCP, and policer parameters
	3.7.0000	Added bits, switch-priority and tc parameters
Example	`switch (config ipv4 access-list my-list)# permit tcp-udp any any eq-destination 100 eq-source 300` `switch (config ipv4 access-list my-list)# permit udp any any eq-destination 100 eq-source 300`	
Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group	
Notes	• It is possible to attach the rule to a unique policer, or to create a policer only for the rule • The policer configuration will always be displayed in bytes • This ACL policer is shared when this table is bound to two or more ports.	

12.12.7.14 deny/permit (IPv4 ICMP ACL rule)

	[seq-number <sequence-number>] {deny | permit} icmp {<source-ip> mask <ip> | any} {<dest-ip> mask <ip> | any} [eq-code <icmp-code>] [eq-type <icmp-type>] [log] [counter | shared-counter <name>] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]} no <sequence-number> Creates a rule for IPv4 ICMP ACL. The no form of the command deletes a rule from the ACL.	
Syntax Description	sequence-number	Optional parameter to set a specific sequence number for the rule Range: 1-65535
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-ip> mask <ip> | any	Sets source IP and optionally sets a mask for that IP address. The “any” option will cause the rule not to check the source IP.
	<dest-ip> mask <ip> | any	Sets destination IP and optionally sets a mask for that IP. The “any” option will cause the rule not to check the destination IP.
	eq-code	Matches ICMP code value. Range: 0-255.
	eq-type	Matches ICMP type value. Range: 0-255.
	log	Enables the log option
	counter	Attaches a unique counter to rule
	shared-counter	Attaches a predefined shared-counter to rule
	ecn	ECN ACL filter. Value: 0-3.
	ttl	Time to live ACL filter. Value: 0-225.
	dscp	DSCP ACL filter. Value: 0-63.
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer. Min value: 8000 bits.
	packets	Attaches packets type policer
	rate	Policer rate value: 100-1000000000000
	k | m | g	Specifies kilo, mega, giga
	burst	Sets burst to policer. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000. For bits there is no default burst. Min value: 2000 bytes.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority. valid values 0-7
	tc <tc_value>	Mapping of matched traffic to tc. valid values 0-7

Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv4 acl	
History	3.1.1400	
	3.5.1000	Updated seq-number parameter
	3.6.2002	Added ICMP parameters
	3.6.5000	Updated command syntax
	3.6.6000	Added ECN, TTL, DSCP, and policer parameters
	3.7.0000	Added bits, switch-priority and tc parameters
Example	`switch (config ipv4 access-list my-list)# permit icmp any any eq-code 10 eq-type 155`	
Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group	
Notes	• ICMP code must be specified in conjunction with an ICMP type. If ICMP type is specified but no ICMP code is specified, the rule matches all ICMP packets of the given type • If no ICMP type or code are specified, the rule matches all ICMP packets from the specified source/destination address • It is possible to attach the rule to a unique policer, or to create a policer only for the rule • The policer configuration will always be displayed in bytes • This ACL policer is shared when this table is bound to two or more ports.	

12.12.7.15 deny/permit (IPv6 ACL rule)

	[seq-number <sequence-number>] {permit | deny} ip {<src-ipv6>/<mask-len> | any} {<dest-ipv6>/<mask-len> | any} [action <action-id>] [log] [counter | shared-counter <name>] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]} no <sequence-number> Creates an IPv6 ACL rule with a specific protocol. The no form of the command deletes a rule from the IPv6 ACL.	
Syntax Description	sequence-number	Optional parameter to set a specific sequence number for the rule Range: 1-65535
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<src-ipv6>/<mask-len> | any	Sets source IP and optionally sets a mask for that IP address. The parameter “any” ignores the source IP.
	<dest-ipv6>/<mask-len> | any	Sets destination IP and optionally sets a mask for that IP. The parameter “any” ignores the destination IP.
	action	Action needs to be defined before attaching to rule
	log	Enables the log option
	counter	Attaches a unique counter to rule
	shared-counter	Attaches a predefined shared-counter to rule

	ecn	ECN ACL filter Range: 0-3
	ttl	Time to live ACL filter Range: 0-225
	dscp	DSCP ACL filter Range: 0-63
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer. Min value: 8000 bits.
	packets	Attaches packets type policer
	rate	Policer rate value Range: 100-1000000000000
	k | m | g	Specifies kilo, mega, giga
	burst	Sets burst to policer. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000. For bits there is no default burst. Min value: 2000 bytes.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority Range: 0-7
	tc <tc_value>	Mapping of matched traffic to TC Range: 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv6 acl	
History	3.6.5000	
	3.6.6000	Added ECN, TTL, DSCP, and policer parameters
	3.7.0000	Added bits, switch-priority and tc parameters
Example	`switch (config ipv6 access-list my-list) # permit ip 2:2::/32 any` `switch (config ipv6 access-list my-list) # permit ip any any policer name`	
Related Commands		
Notes	• IPv6 address format is as follows: <A:B:C:D:E:F:G:H>/mask_len • The fields eq-code (icmp-code) and eq-type (eq-type) are valid only for ICMP rules • It is possible to attach the rule to a unique policer, or to create a policer only for the rule • The policer configuration will always be displayed in bytes • This ACL policer is shared when this table is bound to two or more ports.	

12.12.7.16 deny/permit (IPv6 TCP ACL rule)

	[seq-number <sequence-number>] {permit | deny} tcp {<source-ipv6> /<mask-len> | any} {<dest-ipv6> /<mask-len> | any} [src-port <src-port> | src-port-range <from> <to>] [dest-port <dest-port> | dest-port-range <from> <to>] [established | [ack {0 | 1}] [urg {0 | 1}] [rst {0 | 1}] [syn {0 | 1}] [fin {0 | 1}] [psh {0 | 1}] [ns {0 | 1}] [ece {0 | 1}] [cwr {0 | 1}]] [log] [counter | shared-counter <name>] [action <action-id>] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]} no <sequence-number> Creates an IPv6 ACL rule with a specific protocol. The no form of the command deletes a rule from the IPv6 ACL.	
Syntax Description	sequence-number	Optional parameter to set a specific sequence number for the rule Range: 1-65535
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-ipv6> /<mask-len> | any	Sets source IP and optionally sets a mask for that IP address. The “any” option will cause the rule not to check the source IP.
	<dest-ipv6> / <mask-len> | any	Sets destination IP and optionally sets a mask for that IP. The “any” option will cause the rule not to check the destination IP.
	src-port	L4 source port Note: User may only choose one of the following options to configure source port: src-port; eq-source
	src-port-range	Sets a range of L4 source ports to match Note: User may configure either a single source port or a range
	dest-port	L4 destination port Note: User may only choose one of the following options to configure destination port: dest-port; eq-destination
	dest-port-range	Sets a range of L4 destination ports to match Note: User may configure either a single destination port or a range
	action	Action needs to be defined before attaching to rule
	established	Matches flows which are in established state (“ack” or “rst” flags are set)
	ack; urg; rst; syn; fin; psh; ns; ece; cwr	Matches flows with specific flag Possible match: 0 or 1
	log	Enables the log option
	counter	Attaches a unique counter to rule
	shared-counter	Attaches a predefined shared-counter to rule
	ecn	ECN ACL filter Range: 0-3
	ttl	Time to live ACL filter Range: 0-225
	dscp	DSCP ACL filter Range: 0-63.

	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer. Min value: 8000 bits.
	packets	Attaches packets type policer
	rate	Policer rate value Range: 100-1000000000000
	k | m | g	Specifies kilo, mega, giga
	burst	Sets burst to policer. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000. For bits there is no default burst. Min value: 2000 bytes.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority Range: 0-7
	tc <tc_value>	Mapping of matched traffic to TC Range: 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv6 acl	
History	3.6.5000	
	3.6.6000	Added ECN, TTL, DSCP, policer, and flag parameters
	3.7.0000	Added bits, switch-priority, and tc parameters
Example	`switch (config ipv6 access-list my-list) # permit tcp any 10:10:12::/48`	
Related Commands		
Notes	• IPv6 address format is as follows: <A:B:C:D:E:F:G:H>/mask_len • It is possible to attach the rule to a unique policer, or to create a policer only for the rule • The policer configuration will always be displayed in bytes • This ACL policer is shared when this table is bound to two or more ports.	

12.12.7.17 deny/permit (IPv6 TCP-UDP/UDP ACL rule)

	[seq-number <sequence-number>] {permit | deny} {tcp-udp | udp} {<source-ipv6> / <mask-len> | any} {<dest-ipv6> /<mask-len> | any} [src-port <src-port> | src-port-range <from> <to>] [dest-port <dest-port> | dest-port-range <from> <to>] [log] [counter | shared-counter <name>] [action <action-id>] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]} no <sequence-number> Creates an IPv6 ACL rule with a specific protocol. The no form of the command deletes a rule from the IPv6 ACL.	
Syntax Description	sequence-number	Optional parameter to set a specific sequence number for the rule Range: 1-65535
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass

	<source-ipv6> /<mask-len> | any	Sets source IP and optionally sets a mask for that IP address. The “any” option will cause the rule not to check the source IP.
	<dest-ipv6> / <mask-len> | any	Sets destination IP and optionally sets a mask for that IP. The “any” option will cause the rule not to check the destination IP.
	src-port	L4 source port Note: User may only choose one of the following options to configure source port: src-port; eq-source
	src-port-range	Sets a range of L4 source ports to match Note: User may configure either a single source port or a range
	dest-port	L4 destination port Note: User may only choose one of the following options to configure destination port: dest-port; eq-destination
	dest-port-range	Sets a range of L4 destination ports to match Note: User may configure either a single destination port or a range
	action	Action needs to be defined before attaching to rule
	log	Enables the log option
	counter	Attaches a unique counter to rule
	shared-counter	Attaches a predefined shared-counter to rule
	ecn	ECN ACL filter Range: 0-3
	ttl	Time to live ACL filter Range: 0-225
	dscp	DSCP ACL filter Range: 0-63.
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer. Min value: 8000 bits.
	packets	Attaches packets type policer
	rate	Policer rate value Range: 100-1000000000000
	k | m | g	Specifies kilo, mega, giga
	burst	Sets burst to policer. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000. For bits there is no default burst. Min value: 2000 bytes.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority Range: 0-7
	tc <tc_value>	Mapping of matched traffic to TC Range: 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	

Configuration Mode	config ipv6 acl	
History	3.6.5000	
	3.6.6000	Added ECN, TTL, DSCP, and policer parameters
	3.7.0000	Added bits, switch-priority and tc parameters
Example	`switch (config ipv6 access-list my-list) # permit udp 2:2::/32 10:10:12::/48`	
Related Commands		
Notes	• IPv6 address format is as follows: <A:B:C:D:E:F:G:H>/mask_len • It is possible to attach the rule to a unique policer, or to create a policer only for the rule • The policer configuration will always be displayed in bytes • This ACL policer is shared when this table is bound to two or more ports.	

12.12.7.18 deny/permit (IPv6 ICMPv6 ACL rule)

	[seq-number <sequence-number>] {permit | deny} icmpv6 {<source-ipv6> /<mask-len> | any} {<dest-ipv6> /<mask-len> | any} [code <icmp-code>] [type <icmp-type>] [log] [counter | shared-counter <name>] [action <action-id>] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]} no <sequence-number> Creates an IPv6 ACL rule with a specific protocol. The no form of the command deletes a rule from the IPv6 ACL.	
Syntax Description	sequence-number	Optional parameter to set a specific sequence number for the rule Range: 1-65535
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-ipv6> / <mask-len> | any	Sets source IP and optionally sets a mask for that IP address. The “any” option will cause the rule not to check the source IP.
	<dest-ipv6> / <mask-len> | any	Sets destination IP and optionally sets a mask for that IP. The “any” option will cause the rule not to check the destination IP.
	eq-code	Matches ICMP code value Range: 0-255
	eq-type	Matches ICMP type value Range: 0-255
	action	Action needs to be defined before attaching to rule
	log	Enables the log option
	counter	Attaches a unique counter to rule
	shared-counter	Attaches a predefined shared-counter to rule
	ecn	ECN ACL filter Range: 0-3
	ttl	Time to live ACL filter Range: 0-225

	dscp	DSCP ACL filter Range: 0-63
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer. Min value: 8000 bits.
	packets	Attaches packets type policer
	rate	Policer rate value Range: 100-1000000000000
	k | m | g	Specifies kilo, mega, giga
	burst	Sets burst to policer. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000. For bits there is no default burst. Min value: 2000 bytes.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority Range: 0-7
	tc <tc_value>	Mapping of matched traffic to TC Range: 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv6 acl	
History	3.6.5000	
	3.6.6000	Added ECN, TTL, DSCP, and policer parameters
	3.7.0000	Added bits, switch-priority, and tc parameters
Example	`switch (config ipv6 access-list my-list) # permit icmpv6 any any eq-code 10 eq-type 155`	
Related Commands		
Notes	• IPv6 address format is as follows: <A:B:C:D:E:F:G:H>/mask_len • It is possible to attach the rule to a unique policer, or to create a policer only for the rule • The policer configuration will always be displayed in bytes • This ACL policer is shared when this table is bound to two or more ports.	

12.12.7.19 deny/permit (MAC UDK ACL rule)

	[seq-number <sequence-number>] {deny | permit} {<source-mac> mask <mac-mask> | any} {<dest-mac> mask <mac-mask> | any} [protocol <protocol-num>] [cos <cos>] [vlan <vlan-id>] [vlan-mask <vlan_mask>] [action <action-name>] [log] [counter | shared-counter <name>] [udk <udk1> <val> [mask <mask>]] [<udk2> <val> [mask <mask>]] [<udk3> <val> [mask <mask>]] [<udk4> <val> [mask <mask>]] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]} no <sequence-number> Creates a MAC-UDK ACL rule. The no form of the command deletes a rule from MAC UDK ACL.

Syntax Description	sequence-number	Optional parameter to set a specific sequence number for the rule Range: 1-65535
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-mac> mask <mac-mask> | any	Sets source MAC and optionally sets a mask for that MAC. The “any” option will cause the rule not to check the source MAC.
	<dest-mac> mask <mac-mask> | any	Sets destination MAC and optionally sets a mask for that MAC. The “any” option will cause the rule not to check the destination MAC.
	protocol	Sets the Ethertype filed value from the MAC address Range: 0x0000-0xffff
	cos	Sets the COS (priority bit) field Range: 0-7
	vlan <vlan-id>	Sets the VLAN ID field Range: 1-4094
	vlan-mask <vlan-mask>	Sets VLAN group Range: 0x0000-0x0FFF
	action	Action name (free string)
	log	Enable the log option
	counter	Attach a unique counter to rule
	shared-counter	Attach a predefined shared-counter to rule
	udk	UDK name must be set by user before the rule configuration
	val	The value of the UDK (up to 4 bytes)
	mask	Mask for the UDK value
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer. Min value: 8000 bits.
	packets	Attaches packets type policer
	rate	Policer rate value Range: 100-1000000000000
	k | m | g	Specifies kilo, mega, giga
	burst	Sets burst to policer. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000. For bits there is no default burst. Min value: 2000 bytes.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority Range: 0-7
	tc <tc_value>	Mapping of matched traffic to TC Range: 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	

Configuration Mode	config mac-udk acl	
History	3.6.5000	
	3.6.6000	Added policer parameters
	3.7.0000	Added bits, switch-priority and tc parameters
Example	`switch (config mac-udk access-list mac_udk_acl) # permit any any udk myUdk 10 mask 0xff`	
Related Commands		
Notes	• User cannot attach a shared counter defined on a different ACL table • The parameter shared-counter must be defined before attaching it to the scope of the ACL table • UDK fields must come at the end of the rule configuration • The default mask is 0xff-0xffffffff (depends on value length) • UDK cannot be deleted while it is attached to a rule • 1-4 UDKs per rule may be configured • Values and masks of the UDK can be decimal or hexadecimal • It is possible to attach the rule to a unique policer, or to create a policer only for the rule • The policer configuration will always be displayed in bytes • This ACL policer is shared when this table is bound to two or more ports.	

12.12.7.20 deny/permit (IPv4 UDK ACL rule)

	[seq-number <sequence-number>] {permit | deny} ip {<source-ip> mask <ip> | any} {<dest-ip> mask <ip> | any} [mask <mask>]] [<udk2> <val> [mask <mask>]] [<udk3> <val> [mask <mask>]] [<udk4> <val> [mask <mask>]] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]} no <sequence-number> Creates a rule for IPv4 ACL. The no form of the command deletes a rule from the IPv4 ACL.	
Syntax Description	sequence-number	Optional parameter to set a specific sequence number for the rule Range: 1-65535
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	{any | <source-ip> mask <ip>}	Sets source IP and optionally sets a mask for that IP address. The “any” option causes the rule to not check the source IP. Range: 0-255.
	{any | <destination-ip> mask <ip>}	Sets destination IP and optionally sets a mask for that IP. The “any” option causes the rule to not check the destination IP.
	action	Action needs to be defined before attaching to rule
	log	Enable the log option
	counter	Attach a unique counter to rule
	shared-counter	Attach a predefined shared-counter to rule
	udk	UDK name must be set by user before the rule configuration
	val	The value of the UDK (up to 4 bytes)

	mask	Mask for the UDK value
	ecn	ECN ACL filter| Range: 0-3
	ttl	Time to live ACL filter Range: 0-225
	dscp	DSCP ACL filter Range: 0-63
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer. Min value: 8000 bits.
	packets	Attaches packets type policer
	rate	Policer rate value Range: 100-1000000000000
	k | m | g	Specifies kilo, mega, giga
	burst	Sets burst to policer. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000. For bits there is no default burst. Min value: 2000 bytes.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority Range: 0-7
	tc <tc_value>	Mapping of matched traffic to TC Range: 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv4 acl	
History	3.6.5000	
	3.6.6000	Added ECN, TTL, DSCP, and policer parameters
	3.7.0000	Added bits, switch-priority and tc parameters
Example	`switch (config ipv4 access-list my-list) # deny ip any any action act shared-counter`	
Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group	
Notes	• User cannot attach a shared counter defined on a different ACL table • The parameter shared-counter must be defined before attaching it to the scope of the ACL table • UDK fields must come at the end of the rule configuration • The default mask is 0xff-0xffffffff (depends on value length) • UDK cannot be deleted while it is attached to a rule • 1-4 UDKs per rule may be configured • Values and masks of the UDK can be decimal or hexadecimal • It is possible to attach the rule to a unique policer, or to create a policer only for the rule • The policer configuration will always be displayed in bytes • This ACL policer is shared when this table is bound to two or more ports.	

12.12.7.21 deny/permit (IPv4 TCP UDK ACL rule)

	[seq-number <sequence-number>] {deny | permit} tcp {<source-ip> mask <ip> | any} {<dest-ip> mask <ip> | any} [src-port <src-port> | eq-source <src-port> | src-port-range <from> <to>] [dest-port <dest-port> | eq-destination <dest-port> | dest-port-range <from> <to>] [action <action-id>] [established | [ack {0 | 1}] [urg {0 | 1}] [rst {0 | 1}] [syn {0 | 1}] [fin {0 | 1}] [psh {0 | 1}] [ns {0 | 1}] [ece {0 | 1}] [cwr {0 | 1}]] [log] [counter | shared-counter <name>] [udk <udk1> <val> [mask <mask>]] [<udk2> <val> [mask <mask>]] [<udk3> <val> [mask <mask>]] [<udk4> <val> [mask <mask>]] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]} no <sequence-number> Creates a rule for IPv4 TCP ACL. The no form of the command deletes a rule from the ACL.	
Syntax Description	sequence-number	Optional parameter to set a specific sequence number for the rule Range: 1-65535
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-ip> [mask <ip>] | any	Sets source IP and optionally sets a mask for that IP address. The “any” option will cause the rule not to check the source IP.
	<dest-ip> [mask <ip>] | any	Sets destination IP and optionally sets a mask for that IP. The “any” option will cause the rule not to check the destination IP.
	src-port	L4 source port Note: User may only choose one of the following options to configure source port: src-port; eq-source
	eq-source <src-port>	TCP source port number Range: 0-65535
	src-port-range	Sets a range of L4 source ports to match Note: User may configure either a single source port or a range
	dest-port	L4 destination port Note: User may only choose one of the following options to configure destination port: dest-port; eq-destination
	eq-destination <dest-port>	TCP destination port number Range: 0-65535
	dest-port-range	Sets a range of L4 destination ports to match Note: User may configure either a single destination port or a range
	action	Action needs to be defined before attaching to rule
	established	Matches flows which are in established state (“ack” or “rst” flags are set)
	ack; urg; rst; syn; fin; psh; ns; ece; cwr	Matches flows with specific flag Possible match: 0 or 1
	log	Enables the log option
	counter	Attaches a unique counter to rule
	shared-counter	Attaches a predefined shared-counter to rule

	udk	UDK name must be set by user before the rule configuration
	val	The value of the UDK (up to 4 bytes)
	mask	Mask for the UDK value
	ecn	ECN ACL filter Range: 0-3
	ttl	Time to live ACL filter Range: 0-225
	dscp	DSCP ACL filter Range: 0-63
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer. Min value: 8000 bits.
	packets	Attaches packets type policer
	rate	Policer rate value Range: 100-1000000000000
	k | m | g	Specifies kilo, mega, giga
	burst	Sets burst to policer. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000. For bits there is no default burst. Min value: 2000 bytes.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority Range: 0-7
	tc <tc_value>	Mapping of matched traffic to TC Range: 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv4 acl	
History	3.6.5000	
	3.6.6000	Added ECN, TTL, DSCP, policer, and flag parameters
	3.7.0000	Added bits, switch-priority and tc parameters
Example	`switch (config ipv4 access-list my-list)# permit tcp any any src-port 200 dest-port-range 200 400 established`	
Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group	
Notes	• UDK fields must come at the end of the rule configuration • The default mask is 0xff-0xffffffff (depends on value length) • UDK cannot be deleted while it is attached to a rule • 1-4 UDKs per rule may be configured • It is possible to attach the rule to a unique policer, or to create a policer only for the rule • The policer configuration will always be displayed in bytes • This ACL policer is shared when this table is bound to two or more ports.	

12.12.7.22 deny/permit (IPv4 TCP-UDP/UDP UDK ACL rule)

	[seq-number <sequence-number>] {deny | permit} {tcp-udp | udp} {<source-ip> mask <ip> | any} {<dest-ip> mask <ip> | any} [src-port <src-port> | eq-source <src-port> | src-port-range <from> <to>] [dest-port <dest-port> | eq-destination <dest-port> | dest-port-range <from> <to>] [action <action-id>] [log] [counter | shared-counter <name>] [udk <udk1> <val> [mask <mask>]] [<udk2> <val> [mask <mask>]] [<udk3> <val> [mask <mask>]] [<udk4> <val> [mask <mask>]] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]} no <sequence-number> Creates a rule for IPv4 TCP-UDP/UDP ACL. The no form of the command deletes a rule from the ACL.	
Syntax Description	sequence-number	Optional parameter to set a specific sequence number for the rule Range: 1-65535
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-ip> mask <ip> | any	Sets source IP and optionally sets a mask for that IP address. The “any” option will cause the rule not to check the source IP.
	<dest-ip> mask <ip> | any	Sets destination IP and optionally sets a mask for that IP. The “any” option will cause the rule not to check the destination IP.
	src-port	L4 source port Note: User may only choose one of the following options to configure source port: src-port; eq-source
	eq-source <src-port>	TCP-UDP/UDP source port number Range: 0-65535
	src-port-range	Sets a range of L4 source ports to match Note: User may configure either a single source port or a range
	dest-port	L4 destination port Note: User may only choose one of the following options to configure destination port: dest-port; eq-destination
	eq-destination <dest-port>	TCP-UDP/UDP destination port number Range: 0-65535
	dest-port-range	Sets a range of L4 destination ports to match. Note: User may configure either a single destination port or a range.
	action	Action needs to be defined before attaching to rule
	log	Enables the log option
	counter	Attaches a unique counter to rule
	shared-counter	Attaches a predefined shared-counter to rule
	udk	UDK name must be set by user before the rule configuration
	val	The value of the UDK (up to 4 bytes)
	mask	Mask for the UDK value
	ecn	ECN ACL filter Range: 0-3

	ttl	Time to live ACL filter Range: 0-225
	dscp	DSCP ACL filter Range: 0-63
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer. Min value: 8000 bits.
	packets	Attaches packets type policer
	rate	Policer rate value Range: 100-1000000000000
	k | m | g	Specifies kilo, mega, giga
	burst	Sets burst to policer. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000. For bits there is no default burst. Min value: 2000 bytes.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority Range: 0-7
	tc <tc_value>	Mapping of matched traffic to TC Range: 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv4 acl	
History	3.6.5000	
	3.6.6000	Added ECN, TTL, DSCP, and policer parameters
	3.7.0000	Added bits, switch-priority and tc parameters
Example	`switch (config ipv4 access-list my-list)# permit tcp-udp any any eq-destination 100 eq-source 300` `switch (config ipv4 access-list my-list)# permit udp any any eq-destination 100 eq-source 300`	
Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group	
Notes	• UDK fields must come at the end of the rule configuration • The default mask is 0xff-0xffffffff (depends on value length) • UDK cannot be deleted while it is attached to a rule • 1-4 UDKs per rule may be configured • It is possible to attach the rule to a unique policer, or to create a policer only for the rule • The policer configuration will always be displayed in bytes • This ACL policer is shared when this table is bound to two or more ports.	

12.12.7.23 deny/permit (IPv4 ICMP UDK ACL rule)

	[seq-number <sequence-number>] {deny | permit} icmp {<source-ip> mask <ip> | any} {<dest-ip> mask <ip> | any} [eq-code <icmp-code>] [eq-type <icmp-type>] [log] [counter | shared-counter <name>] [udk <udk1> <val> [mask <mask>]] [<udk2> <val> [mask <mask>]] [<udk3> <val> [mask <mask>]] [<udk4> <val> [mask <mask>]] [ecn <val>] [ttl <val>] [dscp <val>] [policer {<name> | [bytes | packets] rate <rate_value> [k | m | g] [burst <burst_value> [k | m | g]]} no <sequence-number> Creates a rule for IPv4 ICMP ACL. The no form of the command deletes a rule from the ACL.	
Syntax Description	sequence-number	Optional parameter to set a specific sequence number for the rule Range: 1-65535
	deny	Drop all matching traffic
	permit	Allow matching traffic to pass
	<source-ip> mask <ip> | any	Sets source IP and optionally sets a mask for that IP address. The “any” option will cause the rule not to check the source IP.
	<dest-ip> mask <ip> | any	Sets destination IP and optionally sets a mask for that IP. The “any” option will cause the rule not to check the destination IP.
	eq-code	Matches ICMP code value Range: 0-255
	eq-type	Matches ICMP type value Range: 0-255
	log	Enables the log option
	counter	Attaches a unique counter to rule
	shared-counter	Attaches a predefined shared-counter to rule
	udk	UDK name must be set by user before the rule configuration
	val	The value of the UDK (up to 4 bytes)
	mask	Mask for the UDK value
	ecn	ECN ACL filter Range: 0-3
	ttl	Time to live ACL filter Range: 0-225
	dscp	DSCP ACL filter Range: 0-63
	policer	Attaches shared policer to a rule
	bytes	Attaches bytes type policer
	bits	Attaches bits type policer. Min value: 8000 bits.
	packets	Attaches packets type policer
	rate	Policer rate value Range: 100-1000000000000
	k | m | g	Specifies kilo, mega, giga

	burst	Sets burst to policer. If no burst is configured, the default value for type “packets” is 100 and for “bytes” is 10000. For bits there is no default burst. Min value: 2000 bytes.
	switch-priority <switch-priority_value>	Mapping of matched traffic to switch-priority Range: 0-7
	tc <tc_value>	Mapping of matched traffic to TC Range: 0-7
Default	No rule is added by default to access control list Default sequence number is by increments of 10	
Configuration Mode	config ipv4 acl	
History	3.6.5000	
	3.6.6000	Added ECN, TTL, DSCP, and policer parameters
	3.7.0000	Added bits, switch-priority and tc parameters
Example	`switch (config ipv4 access-list my-list)# permit icmp any any eq-code 10 eq-type 155`	
Related Commands	{ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group	
Notes	• ICMP code must be specified in conjunction with an ICMP type. If ICMP type is specified but no ICMP code is specified, the rule matches all ICMP packets of the given type. • If no ICMP type or code are specified, the rule matches all ICMP packets from the specified source/destination address. • UDK fields must come at the end of the rule configuration • The default mask is 0xff-0xffffffff (depends on value length) • UDK cannot be deleted while it is attached to a rule • 1-4 UDKs per rule may be configured • It is possible to attach the rule to a unique policer, or to create a policer only for the rule • The policer configuration will always be displayed in bytes • This ACL policer is shared when this table is bound to two or more ports.	

12.12.7.24 port access-group (IPv4/IPv4 UDK/IPv6/MAC/MAC UDK)

	{ipv4 | ipv4-udk | ipv6 | mac | mac-udk} port access-group <acl-name> no {mac | ipv4 | ipv6 | mac-udk | ipv4-udk} port access-group Attaches an ACL table with bind-point RIF to a VLAN interface. The no form of the command unmaps ACL table with bind-point RIF from a VLAN interface.	
Syntax Description	acl-name	ACL table name
Default	N/A	
Configuration Mode	config interface vlan	
History	3.6.5000	
Example	`switch (config interface vlan 10)# ipv4 port access-group ipv4_acl2`	
Related Commands	show access list summary	

Notes	• Only ACL tables with bind-point set to RIF can be attached to a VLAN interface • Interface VLAN must be configured before binding operation

12.12.7.25 access-list action

	access-list action <action-profile-name> no access-list action <action-profile-name> Creates access-list action profile and entering the action profile configuration mode. The no form of the command deletes the action profile.	
Syntax Description	action-profile-name	Given name for the profile
Default	N/A	
Configuration Mode	config	
History	3.2.0230	
Example	`switch (config)# access-list action my-action` `switch (config access-list action my-action)#`	
Related Commands		
Notes		

12.12.7.26 access-list log

	access-list log [interval <int_num>] [memory <packet_num>] [syslog <packet_num>] no access-list log [interval <int_num>] [memory <packet_num>] [syslog <packet_num>] Configures access list logger. The no form of the command resets parameters for access list logger.	
Syntax Description	interval	Logging interval length in minutes Range: 1min-24hrs
	memory	Maximal number of packets to save in memory Range: 1-3600
	syslog	Maximal number of packets to show in syslog Range: 1-3600
Default	N/A	
Configuration Mode	config	
History	3.6.5000	

Example	`switch (config)# access-list log interval 10`
	`switch (config)# access-list log memory 300`
	`switch (config)# access-list log syslog 200`	

Related Commands		
Notes	• The packet number in syslog configuration must not be greater than the maximal packets number in memory • When configuring interval, the interval will restart resulting in a log dump to syslog and memory clear	

12.12.7.27 vlan-map

	vlan-map <vid> no vlan-map Adds action to map a new VLAN to the packet (in the ingress port or VLAN). The no form of the command removes the action to map a new VLAN.	
Syntax Description	vid	VLAN ID Range: 1-4094
Default	N/A	
Configuration Mode	config acl action	
History	3.2.0230	
Example	`switch (config access-list action my-action)# vlan-map 10`	
Related Commands		
Notes		

12.12.7.28 vlan-pop

	vlan-pop Pops VLAN frames from traffic.
Syntax Description	N/A
Default	N/A
Configuration Mode	config acl action
History	3.4.3000
Example	`switch (config access-list action my-action)# vlan-pop`
Related Commands	
Notes	

12.12.7.29 vlan-push

	vlan-push <vid> Pushes (or adds) VLAN frames to traffic.	
Syntax Description	vid	VLAN ID Range: 1-4094
Default	N/A	
Configuration Mode	config acl action	
History	3.4.3000	
Example	`switch (config access-list action my-action)# vlan-push 10`	
Related Commands		
Notes		

12.12.7.30 monitor session

	monitor session <session_id> Mirrors traffic to monitor session.	
Syntax Description	session_id	The monitor session. Range: 1-3
Default	N/A	
Configuration Mode	config acl action	
History	3.9.3100	
Example	`switch (config access-list action my-action)# monitor session 1`	
Related Commands		

12.12.7.31 show ipv4 access-lists

	show ipv4 access-lists <access-list-name> Displays configuration of IPv4 rules in a specific table.	
Syntax Description	access-list-name	ACL name
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.1400	
	3.3.4500	Updated example
	3.6.6000	Updated example

Example

```
switch (config) # show ipv4 access-lists my-list

Table Type: ipv4
Table Name: my-list
Bind-point: port

------------------------------------------------------------------------------------------------------------------------
seq-number  p/d     protocol  s-ipv4            d-ipv4            sport/type   end-sport  dport/code
end-dport  tcp-control         action  counter     Packets       ttl   ecn   dscp  policer     log
------------------------------------------------------------------------------------------------------------------------
10          permit  ip        any               any               any          none       any
none       N/A                 none    N/A         N/A           none  none  none  none        NO
20          permit  ip        any               any               any          none       any
none       N/A                 none    N/A         N/A           none  none  none  YES         NO
```

Related Commands	deny/permit {ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group
Notes	

12.12.7.32 show ipv4-udk access-lists

	show ipv4-udk access-lists <access-list-name> Displays configuration of IPv4 UDK rules in a specific table.	
Syntax Description	access-list-name	ACL name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
	3.6.6000	Updated example
Example		

```
switch (config) # show ipv4-udk access-lists my-list

Table Type: ipv4-udk
Table Name: my-list
Bind-point: port

-------------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------------------------------------
---------------
seq-number  p/d     protocol  s-ipv4           d-ipv4            sport/type   end-sport  dport/code   end-
dport  tcp-control         action  counter     Packets       udk                            ttl   ecn   dscp
policer    log
-------------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------------------------------------
---------------
7           permit  tcp       any              any               any          none       any          none
any                 none    N/A         N/A                                        none  none  none  none
NO
8           deny    tcp       1.1.1.1/32       any               any          none       any          none
-U        +F        none    N/A         N/A          aaa value 5                  none  none  none  none
NO
10          permit  tcp       1.1.1.1/32       2.2.2.2/32        any          none       any          none
+P-R                none    N/A         N/A          bbb value 6 mask 0x8         none  none  none  none
NO
```

Related Commands	deny/permit {ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group
Notes	

12.12.7.33 show ipv6 access-lists

	show ipv6 access-lists <access-list-name> Displays configuration of IPv6 rules in a specific table.	
Syntax Description	access-list-name	ACL name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
	3.6.6000	Updated example
Example		

```
switch (config) # show ipv6 access-lists my-list

Table Type: ipv6
Table Name: my-list
Bind-point: port

-------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------------
seq-number  p/d     protocol  s-ipv6            d-ipv6             sport/type   end-sport  dport/code
end-dport  tcp-control         action  counter     Packets       ttl   ecn   dscp  policer     log
-------------------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------------
10          permit  ip        any               any                any          none       any
none       N/A                  none    N/A         N/A           33    none  none  none        YES
20          permit  ip        any               any                any          none       any
none       N/A                  none    N/A         N/A           none  none  none  none        NO
30          permit  ip        any               any                any          none       any
none       N/A                  none    N/A         N/A           none  none  none  none        NO
```

Related Commands	deny/permit {ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group
Notes	

12.12.7.34 show mac access-lists

	show mac access-lists <access-list-name> Displays configuration of MAC rules in a specific table.	
Syntax Description	access-list-name	ACL name
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.1400	
	3.3.4500	Updated example
	3.6.6000	Updated example

Example

```
switch (config) # show mac access-lists my-list

Table Type: mac
Table Name: my-list
Bind-point: port

-------------------------------------------------------------------------------------------------------------
----------------------------------------
seq-number  p/d     smac                dmac               protocol  cos    vlan   vlan-mask  action
counter     Packets       policer     log
-------------------------------------------------------------------------------------------------------------
----------------------------------------
10          permit  any                 any                any       any    any    N/A        none
N/A         N/A           roe         NO
```

Related Commands	deny/permit {ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group
Notes	

12.12.7.35 show mac access-lists summary

	show mac access-lists <access-list-name> Displays configuration of MAC rules in a specific table.	
Syntax Description	access-list-name	ACL name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8100	
Example		

```
switch (config) # show mac access-lists summary

-------------------------------------------------------------------------------------
Table type    Table Name       Bind Point    Total entries    Bound to interfaces
-------------------------------------------------------------------------------------
mac           mac1             port          1                Eth1/16
```

Related Commands	deny/permit {ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group
Notes	

12.12.7.36 show mac-udk access-lists

	show mac-udk access-lists <access-list-name> Displays configuration of MAC UDK rules in a specific table.	
Syntax Description	access-list-name	ACL name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
	3.6.6000	Updated example
Example		

```
switch (config) # show mac-udk access-lists my-list

Table Type: mac
Table Name: my-list
Bind-point: port

-------------------------------------------------------------------------------------------------------------------
------------------------
seq-number  p/d     smac      dmac     protocol  cos    vlan   vlan-mask  action  counter     Packets
udk     policer     log
-------------------------------------------------------------------------------------------------------------------
------------------------
10          permit  any       any      any       any    any    N/A        none    N/A         0
YES         NO
20          permit  any       any      any       any    any    N/A        none    N/A         N/A
none        NO
```

Related Commands	deny/permit {ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group
Notes	

12.12.7.37 show access-lists action

	show access-lists action <action-profile-name> Displays the access-list action profiles summary.	
Syntax Description	action-profile-name	Filter the table according to the action profile name
	summary	Display summary of the action list
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.0230	
	3.7.1000	Updated example
	3.9.3100	Updated example to reflect ACL-based monitoring
Example	switch (config)# show access-lists action test_action_1 Access-list Action test_action: ---------- Type Mapped_Vlan_ID Mapped_port Counter_set Policer_ID ---------- vlan-map 1 N/A N/A N/A switch (config)# show access-lists action test_action_2 Access-list Action test_action: ---------- Type Monitor_Sesion Mapped_port Counter_set Policer_ID ---------- monitor 1 N/A N/A N/A	
Related Commands		
Notes		

12.12.7.38 show mac-udk access-lists

	show mac-udk access-lists <access-list-name> Displays configuration of MAC UDK rules in a specific table.	
Syntax Description	access-list-name	ACL name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
	3.6.6000	Updated example
Example		

```
switch (config) # show mac-udk access-lists my-list

Table Type: mac
Table Name: my-list
Bind-point: port

-------------------------------------------------------------------------------------------------------------
----------------------
seq-number  p/d     smac       dmac    protocol  cos    vlan   vlan-mask  action  counter     Packets
udk    policer     log
-------------------------------------------------------------------------------------------------------------
----------------------
10          permit  any        any     any       any    any    N/A        none    N/A         0
YES         NO
20          permit  any        any     any       any    any    N/A        none    N/A         N/A
none        NO
```

Related Commands	deny/permit {ipv4/ipv4-udk/ipv6/mac/mac-udk} access-list {ipv4/ipv4-udk/ipv6/mac/mac-udk} port access-group
Notes	

12.12.7.39 show access-lists log config

	show access-lists log config <action-profile-name> Displays the access-list log configuration information.	
Syntax Description	action-profile-name	Filter the table according to the action profile name
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.0230	
	3.6.8008	Updated example

Example	`switch (config)# show access-lists log config`
	
	`access-list log configuration:`
	`  Memory packets    : 1000`
	`  Syslog packets    : 10`
	`  Interval (minutes): 1`	

Related Commands		
Notes		

12.12.7.40 show access-lists policers (ipv4/ipv4-udk/ipv6/mac/mac-udk)

	show {ipv4 | ipv4-udk | ipv6 | mac | mac-udk} access-lists <access-list-name> policers [name | seq-number] Displays all configured policers on a specific ACL table.	
Syntax Description	access-list-name	ACL name
	name	Policer name filter
	seq-number	Filter by sequence number
Default	N/A	

Configuration Mode	Any command mode
History	3.6.5000
Example	

```
switch (config) # show ipv6 access-lists my-list policers

-----------------------------------------------------------------
Name            Type      Rate           Burst     Sequence Number
-----------------------------------------------------------------
pol             packets   1000           200       50,60,70
rom             packets   1000           200       80
N/A             bytes     12345          20000     40
```

Related Commands	
Notes	

12.12.7.41 show access-lists shared-counters (ipv4/ipv4-udk/ipv6/mac/mac-udk)

	show {ipv4 | ipv4-udk | ipv6 | mac | mac-udk} access-lists <access-list-name> shared-counters Displays all configured shared-counters on a specific ACL table.	
Syntax Description	access-list-name	ACL name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	
Example		

```
switch (config mac access-list my-list) # show mac access-lists mac_acl shared-counters

-----------------------------------------------
counter     packets     total Rules   rule IDs
-----------------------------------------------
cnt1        0           3             20 30 40
cnt2        0           2             50 60
cnt3        0           1             70
```

Related Commands	
Notes	• For each configured shared counter it also displays the counter value (packets), the number of rules attached to this counter and the rule IDs • Up to 5 rule IDs are displayed even though there is no limitation on how many rules can be attached to a counter

12.12.7.42 show access-lists summary

	show [ipv4 | mac | ipv6 | ipv4-udk | mac-udk] access-lists summary Displays the summary of number of rules per ACL, and the interfaces attached.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.1400	
	3.6.5000	Updated example

Example
switch (config) # show access-lists summary --- Table type Table Name Bind type Total entries Bound to interfaces --- mac aaa port 0 Mpo55 ipv4 ddd port 1 Eth1/3, Po1 ipv4 ggg rif 0 VlanIf555 ipv6 table1 port 9 Eth1/9

Related Commands	
Notes	

12.12.7.43 show access-lists log

	show access-lists log [last <num>] Displays captured packets on all access list rules.	
Syntax Description	num	Number of packets to show
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	

Example

```
switch (config) # show access-lists log
Log status: Normal

Log MAC rules:
-------------------------------------------------------------------------------
IF    Table(rule)       Source MAC        Dest MAC          Ethertype VLAN  Hits
-------------------------------------------------------------------------------
1/2   mac_al_log(10)    44:44:44:44:44:44 22:22:22:22:22:22 IPv4      N/A   5

Log IPv4 rules:
-------------------------------------------------------------------------------
IF    Table(rule)       Source IPv4       Dest IPv4         Protocol Source Dest Hits
                                                                     port   port
-------------------------------------------------------------------------------
1/3   ipv4_al_lo(10)    1.1.1.1           2.2.2.2           UDP      44     33   11
```

Related Commands	
Notes	

12.12.7.44 show access-lists log config

	show access-lists log config Displays configuration of access-list logger.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.5000

Example	switch (config) # show access-lists log config access-list log configuration: Memory packets: 1000 Syslog packets: 10 Interval (minutes): 60
Related Commands	
Notes	

12.13 User Defined Keys

User defined keys (UDKs) allow defining custom byte keys—that is, groups of bytes that can be matched to a predefined point in the packet (an extraction point, e.g. the start of a MAC header, or an IP header)—which is useful when wanting to make a match with a part of the packet which does not have a dedicated key.

> The maximum number of UDKs is 4.

An extraction point may be defined for each packet type in a UDK. For each extraction point, an offset (from the beginning of the extraction) is defined.

To be able to modify a UDK after attaching it to an ACL rule, it is first necessary to un-match the UDK from the ACL, and then change the match mode of the UDK to none using the command “no udk match mode”.

> Defining a UDK affects the throughput for packets equal or smaller than 128 bytes.

12.13.1 Configuring UDK

To set UDK with ACL on a specific field:

1. Define new user defined key called ipv4_udk. Run:

```
switch (config) # udk ipv4_udk
switch (config udk ipv4_udk) # exit
```

2. Set user defined key ipv4_udk to match on IPV4 header in offset 4 bytes from start of header. Run:

```
switch (config) # udk ipv4_udk extraction point mode l3 packet type ipv4 extraction point start-of-header
offset 4
```

3. Set the len (in bytes) of the field to match on. Run:

```
switch (config) # udk ipv4_udk len 2
```

4. Set the user defined key to work with access list. Run:

```
switch (config) # udk ipv4_udk match mode acl
```

5. Define new access list table called my_acl_table. Run:

```
switch (config) # ipv4-udk access-list my_acl_table
```

6. Set new rule on the access list table with the previously defined user defined key to match 0x1234. Run:

```
switch (config) # ipv4-udk access-list my_acl_table permit ip any any udk ipv4_udk 0x1234
```

7. Bind the access list table to an ethernet interface. Run:

```
switch (config) # interface ethernet 1/1 ipv4-udk port access-group my_acl_table
```

12.13.2 UDK Commands

12.13.2.1 udk

	udk <udk-name> no udk <udk-name> Creates user defined key. The no form of the command deletes user defined key.	
Syntax Description	udk-name	String
Default	N/A	
Configuration Mode	config	
History	3.6.5000	
Example	`switch (config)# udk udk_name` `switch (config udk udk_name)#`	
Related Commands		
Notes	Defining UDK affects the throughput for packets equal or smaller than 128 bytes.	

12.13.2.2 match mode

	match mode <match-mode> no match mode Configures user defined key match mode. The no form of the command resets this parameter to its default.	
Syntax Description	match-mode	Possible values: • acl • all • ecmp
Default	None	
Configuration Mode	config udk	
History	3.6.5000	

Example	`switch (config udk udk_name)# match mode all`
Related Commands	udk <udk-name>
Notes	

12.13.2.3 extraction point

	extraction point mode <mode> [packet type <type> [extraction point <point> [offset <offset>]]] Configures user-defined key extraction point mode.	
Syntax Description	mode	Possible values: • l2 • l3 • l4
	packet type	Sets user defined key packet type. Possible values: • For L2: l2 • For L3: arp; ipv4; ipv6 • For L4: udp
	extraction point	Sets user defined key extraction point. Possible values for: • l2: l2-ether-type; start-of-header • arp: start-of-header • ipv4; ipv6: start-of-header; start-of-payload • udp: start-of-payload
	offset	Sets user defined key extraction point offset Range: 0-126 (even values)
Default	Mode: l3 Default extraction point per packet type: L2: start-of-header ARP; IPv4; IPv6: start-of-header UDP: start-of-payload Offset: 0	
Configuration Mode	config udk	
History	3.6.5000	
Example	`switch (config udk udk_name)# extraction point mode l3 packet type ipv4 extraction point start-of-header offset 2`	
Related Commands	udk <udk-name>	
Notes		

12.13.2.4 len

	len <length> Configures user-defined key length.	
Syntax Description	length	Range: 1-4
Default	4	
Configuration Mode	config udk	
History	3.6.5000	
Example	`switch (config udk udk_name)# len 4`	

Related Commands	udk <udk-name>
Notes	

12.13.2.5 show udk

	show udk [<udk-name>] Displays summary for user-defined keys.	
Syntax Description	udk-name	Displays information about specific UDK
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.5000	

Example	`switch (config)# show udk`
	
	`UDK name: udk_name`
	`Match mode: none`
	`Length: 4`
	`Extraction mode: 13`
	`IPv4 extraction point: start-of-header`
	`IPv4 offset: 22`
	`IPv6 extraction point: start-of-header`
	`IPv6 offset: 0`
	`ARP extraction point: start-of-header`
	`ARP offset: 0`	

Related Commands	udk <udk-name>	
Notes		

12.14 OpenFlow

NVIDIA Onyx supports OpenFlow 1.3. OpenFlow is a network protocol that facilitates direct communication between network systems via Ethernet. Software Defined Networks (SDN) allows a centralist management of network equipment. OpenFlow allows the SDN controller to manage SDN equipment. The OpenFlow protocol allows communication between the OpenFlow controller and OpenFlow agent.

OpenFlow is useful to manage switches and allow applications running on the OpenFlow controller to have access to the switch's data path and provide functionality such as flow steering, security enhancement, traffic monitoring and more.

The OpenFlow controller communicates with the OpenFlow switch over secured channel using OpenFlow protocol.

An OpenFlow switch contains a flow table which contains flows inserted by the OpenFlow controller. And the OpenFlow switch performs packet lookup and forwarding according to those rules.

OpenFlow switch implementation is based on the hybrid model, allowing the coexistence of an OpenFlow pipeline and a normal pipeline. In this model, a packet is forwarded according to OpenFlow configuration, if such configuration is matched with the packet parameters, otherwise the packet is handled by the normal (regular forwarding/routing) pipeline. NVIDIA Onyx allows configuring regular switch port and port-channel ports to be in hybrid mode (this is relevant to regular switch ports, port-channel switch ports, regular router ports and port-channel router ports).

The OpenFlow specification defines:

> *"OpenFlow-hybrid switches support both OpenFlow operation and normal Ethernet switching operation, i.e. traditional L2 Ethernet switching, VLAN isolation, L3 routing (IPv4 routing, IPv6 routing...), ACL and QoS processing. Those switches must provide a classification mechanism outside of OpenFlow that routes traffic to either the OpenFlow pipeline or the normal pipeline. For example, a switch may use the VLAN tag or input port of the packet to decide whether to process the packet using one pipeline or the other, or it may direct all packets to the OpenFlow pipeline."*

Utilizing the built-in capabilities of the hybrid switch/router is the main benefit of the hybrid mode. It increases network performance and efficiency – faster processing of new flows as well as lower load on the controllers. The hybrid switch processes non-OpenFlow data through its local management plane and achieve better efficiency and use of resources, compared to the pure OpenFlow switch.

- Flow Table
- OpenFlow 1.3 Workflow
- Configuring OpenFlow
- Configuring Flows Using CLI Commands
- Configuring Secure Connection to OpenFlow
- OpenFlow Commands

12.14.1 Flow Table

The flow table contains flows which are used to perform packet lookup, modification and forwarding. Each flow has a 12 tuple key. The key is used in order to classify a packet into a certain flow. The key contains the flowing fields: ingress port, source MAC, destination MAC, EtherType, VLAN ID, PCP, source IP, destination IP, IP protocol, IP ToS bits, TCP/UDP source port and TCP/UDP destination port.

The flow key can have a specific value for each field or wildcard which signals to the switch to ignore this part of the key.

Each packet passes through the flow table once a match is found; the switch performs the actions configured to the specific flow by the OpenFlow controller.

Up-keeping a flow table enables the switch to forward incoming traffic with a simple lookup on its flow table entries. OpenFlow switches perform a check for matching entries on, or ignore using a wildcard, specific fields of the ingress traffic. If the entry exists, the switch performs the action associated with that flow entry. Packets without a flow entry match are forwarded according to the normal pipeline (hybrid switch).

Every flow entry contains one of the following parameters:

1. Header fields for matching purposes with each entry containing a specific value or a wildcard which could match all entries.
2. Matching packet counters which are useful for statistical purposes, in order to keep track of the number of packets.
3. Actions which specify the manner in which to handle the packets of a flow which can be any of the following:
   - Forwarding the packet
   - Dropping the packet
   - Forwarding the packet to the OpenFlow controller
   - Modifying the VLAN, VLAN priority (PCP), and/or stripping the VLAN header

12.14.2 OpenFlow 1.3 Workflow

The OpenFlow (OF) pipeline is deployed in parallel to the usual NVIDIA Onyx pipeline.

The ingress port must be deployed in hybrid mode so as to serve both the OF and normal NVIDIA Onyx pipeline.

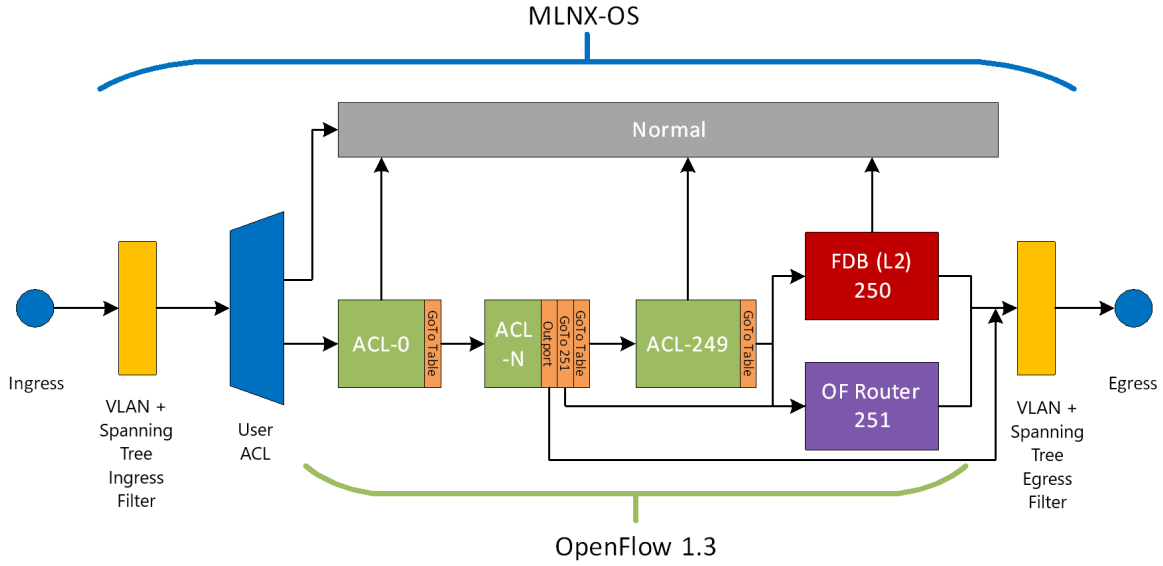

The ingress packet, which passes the VLAN and Spanning Tree filters and is a match to the user ACL table, either progresses to the regular NVIDIA Onyx flow or the OpenFlow pipeline depending on the port coupling.

The following table presents a general summary of the capabilities of the OpenFlow 1.3 pipeline. They are also described further on in the document.

Table	Match	Actions	Group	Meters

ACLs [0-249]	• in_port
	• dl_src
	• dl_dst
	• dl_type
	• vlan_vid
	• vlan_pcp
	• ip_src
	• ip_dst
	• ipv6_dst
	• ipv6_src
	• ip_proto
	• ip_dscp
	• ip_ecn
	• ip_ttl
	• 14_src_port
	• 14_dst_port
	• tunnel_id
	• metadata 0xFFF
	• mpls_label
	• Table must be configured using “openflow table match-keys” to support the following fields:
	  • ip_src_inner
	  • ip_dst_inner
	  • ignr_eth_type
	(Dynamic key)
	(Arbitrary mask)
		• Push/pop VLAN
		• SET_TTL
		• DEC_TTL
		• goto_table
		• Set queue
		• Eth SRC/DST MAC
		• VLAN ID
		• PCP
		• DSCP
		• ECN
		• Output
		• Group
		• Meters
		• Normal
		• controller (trap)
		• L4 src/dst
		• write_metadata
			• ALL—Output ports; Set field
			• Select—{weights} Output ports (without LAG)
			• FF—Output ports
				• KBps/PKTs—{Burst}
				• Drop

FDB [250]	• vlan_vid • dl_dst **(Exact match)**	• OUTPUT • DROP • Normal	Select—{Weights} Output ports (without LAG)	N/A
Router [251]	• ipv4_dst • ipv6_dst **(LPM)**	• DEC_TTL • SET_DMAC • OUTPUT • DROP **(Must have DEC_TTL and SET_DMAC when output action is implemented)**		N/A

Onyx only supports up to 50 actions per flow/group.

12.14.2.1 ACL Rule Tables (0-249)

An Access Control List (ACL) is a list of permissions attached to an object, to filter or match switches packets. When the pattern is matched at the hardware lookup engine, a specified action (e.g. permit/deny) is applied. The rule fields represent flow characteristics such as source and destination addresses, protocol and VLAN ID.

ACL support currently allows actions of permit or deny rules, and supports only ingress direction. ACL search pattern can be taken from either L2 or L3 fields.

12.14.2.1.1 Supported ACL Matching Rules

Ingress packets, arriving the ACL, are matched against any combination of the following parameters (defined as the key):

- OXM_OF_METADATA—matches according to metadata
- OXM_OF_IN_PORT—matches according to ingress port (exact match or wildcard)
- OXM_OF_ETH_SRC—matches source MAC address
- OXM_OF_ETH_DST—matches destination MAC address
- OXM_OF_ETH_TYPE—matches EtherType

  When match rule is set to match eth_type 9100, VLAN ID matching does not work.

- OXM_OF_VLAN_VID—matches VLAN ID
- OXM_OF_VLAN_PCP—matches priority level
- OXM_OF_IPV4_SRC—matches source IPv4 address
- OXM_OF_IPV4_DST—matches destination IPv4 address
- OXM_OF_IPV6_SRC—matches source IPv6 address
- OXM_OF_IPV6_DST—matches destination IPv6 address
- OXM_OF_IPV6_ND_TARGET

  OXM_OF_IPV6_ND_TARGET match rule is not supported.

- OXM_OF_IP_PROTO—matches IP protocols (exact match or wildcard)
- OXM_OF_IP_DSCP—matches IP DSCP field (exact match or wildcard)
- OXM_OF_IP_ECN—matches network ECN (exact match or wildcard)
- OXM_OF_NW_TTL—matches network TTL (exact match or wildcard)
- OXM_OF_TCP_SRC—matches source TCP
- OXM_OF_TCP_DST—matches destination TCP
- OXM_OF_UDP_SRC—matches source UDP
- OXM_OF_UDP_DST—matches destination UDP
- OXM_OF_SCTP_SRC—matches source SCTP
- OXM_OF_SCTP_DST—matches destination SCTP
- OXM_OF_ICMPV4_TYPE—matches ICMP type
- OXM_OF_ICMPV4_CODE—matches ICMP code
- OXM_OF_ARP_OP—matches ARP OP code
- OXM_OF_ARP_SPA—matches sender protocol address
- OXM_OF_ARP_TPA—matches target protocol address

There is a default set of match keys configured. To see what it is, please run the command “show openflow table match-keys” on your machine. To alter it, please use the command “openflow table match-keys”.

12.14.2.1.2 Non-standard Matches

OpenFlow 1.3 is able to match non-standard OpenFlow matching rules by mapping them to standard ones. The following non-standard matches are supported:

- Matching source/destination IPv4 address encapsulated with MPLS labels (up to 6 MPLS labels can be skipped)—ip_src_inner/ip_dst_inner is mapped to OXM_OF_IPV4_SRC, OXM_OF_IPV4_DST
- Table configuration:

```
openflow table 0 match-keys dl_dst dl_src dl_type mpls_label vlan_vid
openflow table 10 match-keys ignr_eth_type ip_dst_inner ip_src_inner
```

  The ignr_eth_type is needed to ignore the Ethertype of IP that is required by OpenFlow to set to as a prerequisite to match on IP addresses.
- Rules:

```
openflow add-flows 1 table=0,mpls,mpls_label:32,actions=goto_table=10
openflow add-flows 2 table=10,ip,nw_src=10.10.10.0/24,nw_dst=10.10.20.0/24,
actions=output:127
```

  The above matches IP address from 10.10.10.0/24 to 10.10.20.0/24 which have MPLS label 32 as the first label.

> Control actions are not supported for non-standard matches.

12.14.2.1.3 Supported Rule Table Instructions

The intercepted packet is processed according to the instructions on the rule tables. The supported instructions are as follows:
DROP—drops packet

- OFPIT_GOTO_TABLE—sends the packet for processing by another rule table
- OFPIT_METER—policer function; drops packet if it exceeds kbps/pktps limit
- OFPIT_WRITE_METADATA—writes meta-data with mask <METADATA>/0xFFF
- OFPIT_EXPERIMENTE—sends the packet for processing by another controller
- OFPIT_APPLY_ACTIONS—applies certain actions specified in the section below

12.14.2.1.4 Supported ACL Apply Actions

The following actions are applied on ingress packets once a match is achieved on the ACL table:

- OFPAT_OUTPUT—the packet is sent out to a port (may also be a controller port)
- OFPAT_GROUP—the packet is sent out to a group
  3 types of group ports are supported:
  - All: The packet is broadcasted on all ports which are part of the defined group
  - Selected: The packets are distributed toward the group ports according to a weight mechanism
  - Fast-Failover (FF): FF is a group of ports, one of which is defined as the primary port through which the packets are transported. In a failure scenario (defined as part of the group definition), traffic becomes transported through the most eligible backup port

(from the list of backup ports). Once the failure scenario ends, traffic is routed again through the primary port

- OFPAT_POP_VLAN—strips 802.1Q (VLAN) tag from the packet
- OFPAT_PUSH_VLAN—adds 802.1Q (VLAN) tag from the packet
- OFPAT_SET_NW_TTL—modifies network TTL
- OFPAT_DEC_NW_TTL—decrements network TTL
- OFPAT_SET_FIELD—ACL set fields detailed in section below
- Normal

12.14.2.1.5 Supported ACL Set Fields

The following modifications may be implemented on ingress packets:

- OXM_OF_ETH_SRC—sets the source MAC address of the packet
- OXM_OF_ETH_DST—sets the destination MAC address of the packet
- OXM_OF_VLAN_VID—sets the VLAN ID of the packet
- OXM_OF_VLAN_PCP—sets the VLAN priority code point (PCP; 0-7)
- OXM_OF_IP_DSCP—sets IP DSCP
- OXM_OF_IP_ECN—sets network ECN
- NXM_NX_CT_NW_SRC*—sets the source IP address of the packet
- NXM_NX_CT_NW_DST*—sets the destination IP address of the packet
- NXM_NX_CT_TP_SRC*—sets the source L4 port of the packet
- NXM_NX_CT_TP_DST*—sets the destination L4 port of the packet

*Supported only on Spectrum-2 and Spectrum-3 systems.

12.14.2.1.6 Supported ACL Meters

- ACL tables support up to 968 meters with 1 band (drop) per meter.
- Valid meter ID range: 1-969
- Only the rate or the burst size fields can be modified using OFPMC_MODIFY
- OFPMF_BURST meter type can be OFPMF_KBPS (KB/s) or OFPMF_PKTPS (number of packets per second) but not both
- Meter actions:
- OFPMBT_DROP—drops packet according to meter configuration

12.14.2.1.7 FDB Table (250)

The FDB table is the same one shared with regular NVIDIA Onyx configuration (e.g., learning, static macs, and so forth). The cumulative number of supported FDB rules is 88KB. FDB may only configure rules with priority of 0x8000. Hard timeout is supported for FDB table rules. FDB rules cannot have wildcard on VID/ETH_DST.

The default action for the FDB table is normal and this cannot be changed by the user.

12.14.2.1.8 Supported FDB Apply Actions

- OFPAT_OUTPUT—the packet is sent out to a port (may be controller port)
- DROP—drops packet
- Normal

12.14.2.1.9 Supported FDB Matching Rules

- OXM_OF_VLAN_VID—matches VLAN ID
- OXM_OF_ETH_DST—matches destination MAC address

12.14.2.2 Router Table (251)

The OpenFlow router table and the regular NVIDIA Onyx router table share the same hardware resources, but are separated logically.

The cumulative number of supported FDB & router rules is 88K. Hard timeout, where the switch removes a rule after a configured timer expires, is supported for router table rules. Switch systems ignore rule priority and configure rules according to masklen in DST IPv4/IPv6 match. A rule with action output must have SET_FIELD with ETH_DST and DEC_NW_TTL. The default action for the router table is DROP.

Set DMAC can be assigned only to one output port. When a new rule with a set DMAC and a new output port is configured, the previous rules are removed from the hardware. Later, if the new configuration is deleted, the previous rules get reinstalled in hardware.
Note that all sent packets from the Router Table are without a VLAN header (untagged).

12.14.2.2.1 Supported Router Apply Actions

- OFPAT_OUTPUT—the packet is sent out to a port (may be controller port)
- OFPAT_DEC_NW_TTL—decrements network TTL
- OFPAT_SET_DMAC—OFPAT_SET_FIELD with OFPXMT_OFB_ETH_DST
- DROP—drops packet

> When an output action is implemented, DEC_TTL and SET_DMAC must also be set.

12.14.2.2.2 Supported Router Set Fields

- OXM_OF_ETH_DST—sets the destination MAC address of the packet

12.14.2.2.3 Supported Router Matching Rules

- OXM_OF_IPV4_DST—matches destination IPv4 address
- OXM_OF_IPV6_DST—matches destination IPv6 address

12.14.3 Configuring OpenFlow

To run OpenFlow on a switch:

1. Unlock the OpenFlow CLI commands. Run:

```
switch (config) # protocol openflow
```

2. Configure interfaces to be managed by OpenFlow. Run:

```
switch (config) # interface ethernet 1/1-1/4 openflow mode hybrid
```

3. Configure the OpenFlow controller IP and TCP port. Run:

```
switch (config) # openflow controller-ip 10.209.0.205 tcp-port 6633
```

12.14.4 Configuring Flows Using CLI Commands

The on-switch commands use the Open vSwitch (OVS) syntax for OpenFlow. They are actually based on the "ovs-ofctl" command. For more details please refer to the Flow Syntax section of this man-page.

It is slightly modified as you need to explicitly input a flow reference number to modify. This flow ID may be used when performing any modification to the flow (e.g. delete).

All flow configurations also appear in the running-config and are restored after switch reload.

When configuring flows, you may assign them a high priority, and then to configure a "drop all" rule for non-matching packets with a lower priority.

For the flows (use a higher priority e.g. 10000 then the drop all rule) and input interface:

```
switch (config) # openflow add-flows 1 ip, priority=5000, in_port=Eth1/1, nw_src=192.168.0.1/32, nw_dst=239.0.1.2/
32, actions=output=Eth1/56
```

The above rule matches on SRC IP=192.168.0.1 and DEST IP=239.0.1.2 and the action is to output matching traffic to interface Eth1/56.

For the "drop all" rule (use a lower priority than other match rules):

```
switch (config) # openflow add-flows 1000 priority=50,in_port=ANY,actions=DROP
```

To delete a flow, run the command "del-flows" along with a flow's reference number:

```
switch (config) # openflow del-flows 1
switch (config) # openflow del-flows 1000
```

> OpenFlow may be configured using one method at a time, so if an OpenFlow controller is configured then switch CLI method cannot be used.

12.14.4.1 Support of MLAG Interface in OpenFlow

To configure MLAG interface in OpenFlow, do the following:

1. Enable OpenFlow in the system.

```
switch (config) # protocol openflow
```

2. Add MPO interfaces as OpenFlow mode hybrid port.

```
switch (config) # interface mlag-port-channel 1-3 openflow mode hybrid
```

3. Add the needed OpenFlow flow with MPO usage.

```
openflow add-flows 1 table=0, priority=500, in_port=Mpo1, actions=NORMAL
```

4. Observe the relevant MPO interfaces in OpenFlow using the "show openflow" command.

```
switch (config) # show openflow
OpenFlow Version: OpenFlow 1.3
Datapath ID: 0000248a07cacd00

Controllers Information:
  ------------------------------------------------------------------------------------------
  Controller                State            Role       Changed (sec)  Last Error
  ------------------------------------------------------------------------------------------


Mapping of OpenFlow ports to their OpenFlow numbers:
  ----------------------
  Interface       OF-Port
  ----------------------
  Mpo1            OF-29001
  Mpo2            OF-29002
  Mpo3            OF-29003
```

5. Observe the OpenFlow rules with MPO interfaces with the "show openflow flows ethernet-names" command.

```
switch (config) # show openflow flows ethernet-names
OFPST_FLOW reply (OF1.3) (xid=0x2):
 cookie=0x0, duration=2.166s, table=0, n_packets=0, n_bytes=0, priority=500,in_port=Mpo1 actions=NORMAL
```

6. Only 63 POs/MPOs interfaces are allowed if protocol OpenFlow is enabled (1 LAG is always used by OpenFlow by default).

```
switch (config) # protocol openflow
switch (config) # protocol mlag
switch (config) # interface mlag-port-channel 1-32
switch (config) # interface port-channel 33-64
 % The one LAG is in use by OpenFlow feature, please disable OpenFlow to have a possibility use 64 POs/
MPOs.
```

7. It is not possible to enable OpenFlow protocol if there are already 64 POs/MPOs in use. Only 63 POs/MPOs can be used as 1 LAG is always used by OpenFlow:

```
switch (config) # protocol mlag
switch (config) # interface mlag-port-channel 1-64
switch (config) # protocol openflow
 % There are already 64 POs/MPOs in use. One free LAG is required to enable protocol OpenFlow.
```

Forwarding control traffic (LACP, LLDP, BPDU) from one MPO interface to another one is currently not supported. If the destination port is MPO and this MPO is in the DOWN state, the traffic will be redirected to IPL and cause unexpected behavior (the IPL will start flapping due to redundant packets, e.g. LACP packets).

For example, it may lead to the scenario when LACP packets will be delivered on the wrong destination port when the following OpenFlow rule will be used:
openflow add-flows 1 table=0,priority=100,in_port=Mpo1,actions=output:Mpo2

The workaround is to use separate OpenFlow rule with a higher priority and destination MAC for the LACP packets in order to forward LACP packets in the proper direction:
openflow add-flows 1 table=0,priority=200,dl_dst= 01:80:c2:00:00:02,actions=NORMAL
openflow add-flows 1 table=0,priority=100,in_port=Mpo1,actions=output:Mpo2

Proper setup of MLAG topology and MLAG failovers are under user responsibility. Openflow does not handle such situations. The OpenFlow rules are also not synchronized on the MLAG members.

12.14.5 Configuring Secure Connection to OpenFlow

Since OpenFlow requires a certificate signed by the certificate authority (CA), the default certificate, which is self-signed, must be replaced.

If using a certificate generated by the switch, skip steps 2 and 3 below.

To change the default certificate for a secure OpenFlow connection:

1. Import the certificate to be used (e.g., a certificate created by openssl outside the switch). Run:

```
switch (config) # crypto certificate name my-openflow public-cert pem "-----BEGIN CERTIFICATE-----
> MIIDYzCCAksCCQC9EPbMuxjNBzANBgkqhkiG9w0BAQsFADBeMQswCQYDVQQGEwJJ
...
> fEt2ui9taB1dl9480xDsGUxwUDX4YOs/bQDjp99z+cKXUe2eYzeEwnTdrCzPZuQo
> -----END CERTIFICATE-----"
Successfully installed certificate with name 'my-openflow'
```

Or use a new self-signed certificate via switch CLI and export it as a CSR (certificate signing request) and send said CSR to the root CA for signing:

```
switch (config) # crypto certificate name my-openflow generate self-signed
Successfully generated certificate with name ' my-openflow'

switch (config) # show crypto certificate name my-openflow csr-pem

-----BEGIN CERTIFICATE REQUEST-----
MIICuDCCAaACAQAwczELMAkGA1UEBhMCSVMxDDAKBgNVBAgMA1RCRDEMMAoGA1UE
BwwDVEJEMQwwCgYDVQQKDANUQkQxDDAKBgNVBAsMA1RCRDEYMBYGA1UEAwwPYnVs
bGRvZy1xcDEtMTMzMRIwEAYJKoZIhvcNAQkBFgNUQkQwggEiMA0GCSqGSIb3DQEB
AQUAA4IBDwAwggEKAoIBAQC34xRVh9BaBUPIilV6kiSOAVAnOFgreWtEYoWeGpWJ
XGZQBwewFx4TGptYo5fZ4KcnYcQxrcW7gYycQB9Y+9vUVvvPi3b4aYc2FkoNtnC3
0BRTxEcIiwXY7LQxIA23Zuv/OlhjTkpe0+OYtpJSFeIDKMIX4Uy2BfevG06YLCAW
tuju2FLQVkexayNK/HFLa5POpVt+16JLB1eV0bcC38Mq9JNIgPspJ7JIjo+BjzgD
43iEY41hlRzoalu78nBBd0HbAddxCF1Uc+8PLuPLCIjGbV9ehPJNWSsA/T9jUEFU
90KaI0/k05JqCXWnpvKz3opQraHsVAbsxG312pnmbTFNAgMBAAGgADANBgkqhkiG
9w0BAQsFAAOCAQEAhpgZRNW/jleyhUbtGEr0CzdNbJ70V8w2lGr6bDhZgrQ/I4eO
1K1D1hvfrVWYRB0SSPFmCmVmFmC7BQne8xrbL2It3ZdSKd82Ts36/Uxjtb63hyt3
GBzCas7qypsbCVW42UHuD+259Yu5xpi9haspzD8Wg2ZKU5e6SjcH+JIchkM9mh/g
BQJo4shybTgPfT+mFUCCygWmf5aLyQ9TrZpaUQ7cOk6BZB1RRkOVvA6uCfrwlBks
X72LleceL4fP9dtML4VMzMMAf+wOUNxWP9+lqkKMaDhroDP5qlo/lr5BLSlRVet4
z7zb3xSaPrhnefoGr88WFO74d9RxLPPdHcfMFw==
-----END CERTIFICATE REQUEST-----
```

2. Import key of certificate. Run:

```
switch (config) # crypto certificate name my-openflow private-key pem "-----BEGIN RSA PRIVATE KEY-----
> MIIEpAIBAAKCAQEAypJnZkwbhmt71Kf/MO6cy7QmWWHhCozzWRwuWGKse+MxSmfC
...
> QAuPOVR1lSyIEnYU+X0rMHc/9tgUh/8C7mBKwj7dccMmnRWz2djsjg==
> -----END RSA PRIVATE KEY-----"
```

3. Designate “my-openflow” as the global default certificate for authentication of this system to clients. Run:

```
switch (config) # crypto certificate default-cert name my-openflow
```

4. Import the CA certificate which signed for the controller. Run:

```
switch (config) # # crypto certificate name rootCA public-cert pem "-----BEGIN CERTIFICATE-----
> MIIDjzCCAnegAwIBAgIJALVou4mcQtxlMA0GCSqGSIb3DQEBCwUAMF4xCzAJBgNV
...
> +ZfQIOCFS8gY4BDq73W4ugr38mqIA8UXXAMPwgjCbk4NyOh0rJ1P6WT8fYzvunct
> -----END CERTIFICATE-----"
Successfully installed certificate with name 'rootCA'
```

5. Adds the “rootCA” to the default CA certificate list. Run:

```
switch (config) # crypto certificate ca-list default-ca-list name rootCA
```

6. Save configuration. Run:

```
switch (config) # configuration write
```

7. Reboot the switch. Run:

```
switch (config) # reload
```

8. Verify configuration. Run:

```
switch (config) # show crypto certificate
Certificate with name 'system-self-signed'
    Comment:                        system-generated self-signed certificate
    Private Key:                    present
    Serial Number:                  0x543e2efc3a5ecdbe18b5b5e744598424
    SHA-1 Fingerprint:              14e1d36035c7a5fea9f7f0f423572c9954cb9fac

    Validity:
        Starts:                     2016/09/12 12:44:10
        Expires:                    2017/09/12 12:44:10
 Subject:
        Common Name:                switch
        Country:                    IS
        State or Province:          TBD
        Locality:                   TBD
        Organization:               TBD
        Organizational Unit:        TBD
        E-mail Address:             TBD

    Issuer:
        Common Name:                switch
        Country:                    IS
        State or Province:          TBD
        Locality:                   TBD
        Organization:               TBD
        Organizational Unit:        TBD
        E-mail Address:             TBD

Certificate with name 'my-openflow' (default-cert)
    Private Key:                    present
    Serial Number:                  0xbd10f6ccbb18cd07
    SHA-1 Fingerprint:              1e0e3302182ab56f2cbd3ca21722dec55299d670

    Validity:
        Starts:                     2016/09/12 15:16:48
        Expires:                    2018/01/25 14:16:48

    Subject:
        Common Name:                switch
        Country:                    *
        State or Province:          Some-State
        Locality:                   *
        Organization:               Mlnx
        Organizational Unit:        e2e
        E-mail Address:             none@nowhere.com

    Issuer:
        Common Name:                ca
        Country:                    *
        State or Province:          Some-State
        Locality:                   *
        Organization:               Mlnx
        Organizational Unit:        e2e
    Certificate with name 'rootCA'
    Private Key:                    not present
    Serial Number:                  0xb568bb899c42dc65
    SHA-1 Fingerprint:              9855536f6ee0177356ffbdc54ffe803bc83fb4c6
    Validity:
        Starts:                     2016/09/08 10:34:23
        Expires:                    2019/06/29 10:34:23

    Subject:
        Common Name:                ca
        Country:                    *
        State or Province:          Some-State
        Locality:                   *
```

```
        Organization:              Mlnx
        Organizational Unit:       e2e

    Issuer:
        Common Name:               ca
        Country:                   *
        State or Province:         Some-State
        Locality:                  *
        Organization:              Mlnx
        Organizational Unit:       e2e
```

9. Configure secure controller IP connection. Run:

```
switch (config) # controller-ip 10.10.10.10 tls
```

12.14.6 OpenFlow Commands

- 12.14.6.1 protocol openflow
- 12.14.6.2 openflow mode hybrid
- 12.14.6.3 openflow add-flows
- 12.14.6.4 openflow del-flows
- 12.14.6.5 openflow add-group
- 12.14.6.6 openflow del-group
- 12.14.6.7 openflow mod-group
- 12.14.6.8 openflow add-meter
- 12.14.6.9 openflow del-meter
- 12.14.6.10 openflow fail-mode secure
- 12.14.6.11 openflow mod-meter
- 12.14.6.12 openflow re-apply flows
- 12.14.6.13 openflow re-apply groups
- 12.14.6.14 openflow re-apply meters
- 12.14.6.15 controller-ip
- 12.14.6.16 datapath-id
- 12.14.6.17 openflow table match-keys
- 12.14.6.18 openflow acl table counter disable
- 12.14.6.19 show openflow
- 12.14.6.20 show openflow flows
- 12.14.6.21 show openflow flows ethernet-names
- 12.14.6.22 show openflow groups
- 12.14.6.23 show openflow groups ethernet-names
- 12.14.6.24 show openflow meters
- 12.14.6.25 show openflow flows table
- 12.14.6.26 show openflow flows cookie
- 12.14.6.27 show openflow table match-keys
- 12.14.6.28 show openflow table match-keys supported

12.14.6.1 protocol openflow

	protocol openflow no protocol openflow Unhides the OpenFlow commands. The no form of the command hides the OpenFlow commands.
Syntax Description	N/A
Default	no protocol openflow
Configuration Mode	config
History	3.3.4200
Example	`switch (config) # protocol openflow`
Related Commands	
Notes	

12.14.6.2 openflow mode hybrid

	openflow mode hybrid no openflow mode Enables OpenFlow on the port. The no form of the command returns the port to its default state.	
Syntax Description	N/A	
Default	no openflow mode	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.3.4200	
	3.6.2100	Updated notes
	3.9.2000	Updated example and notes
	3.9.2400	Added mlag-port-channel configuration option and updated note

Example	`switch (config interface ethernet 1/1)# openflow mode hybrid`
	`switch (config interface port-channel 1)# openflow mode hybrid`
	`switch (config interface mlag-port-channel 1)# openflow mode hybrid`	

Related Commands	protocol openflow	
Notes	Running OpenFlow on the MLAG interface has several limitations. Please see them in the section "Support of MLAG Interface in OpenFlow" section	

12.14.6.3 openflow add-flows

	openflow add-flows <flow-id> [[table-id],[priority-id],<match-parameter1> [,...,<match-parameterN>],<action1>[,...,<actionN>]] Adds OpenFlow flow.	
Syntax Description	flow-id	ID number to give this flow Range: 0-65535

	priority-id	Priority to give this flow Range: 0-65535
	match-parameter	Rule according to which a match is made. For a list of supported matches, see the match column in the "OpenFlow 1.3 Pipeline Capabilities Summary Table".
	table-id	Range: • ACLs: 0-249 • FDB: 250 • Router: 251
	action	Action to perform on the matched traffic. For a list of supported actions, see the action column in "OpenFlow 1.3 Pipeline Capabilities Summary Table".
Default	table-id default is 0	
Configuration Mode	config	
History	3.6.4006	

Example

```
switch (config)# openflow add-flows 1,
priority=10,in_port=Eth1/1,nw_src=192.168.0.1/32,nw_dst=239.0.1.2/32,actions=output=Eth 1/11,Eth 1/22,Eth
1/33

switch (config)# openflow add-flows 3 table=3,in_port=121,actions=output:117

switch (config)# openflow add-flows 2 in_port=ANY,actions=push_vlan:33024,mod_vlan_vid:4111

switch (config)# openflow add-flows 4
table=0,priority=101,dl_type=0x0800,in_port=79,dl_vlan=233,nw_dst=172.0.0.0/8,actions=pop_vlan,goto_table:
251

switch (config)# openflow add-flows 5 in_port=1,actions=dec_ttl

switch (config)# openflow add-flows 6
table=0,priority=777,in_port=121,dl_type=0x0800,nw_proto=6,actions=mod_nw_ttl:55,output:99

switch (config)# openflow add-flows 7
table=0,priority=777,in_port=121,dl_type=0x0800,nw_proto=6,actions=Set_field:55-\>nw_ttl,output:99

switch (config)# openflow add-flows 8
table=0,priority=777,in_port=121,actions=output:99,Set_field:11:22:33:44:00:00-\>eth_dst

switch (config)# openflow add-flows 9
table=0,priority=777,in_port=121,dl_type=0x0800,nw_proto=6,actions=Set_field:0-\>ip_ecn,output:99

switch (config)# openflow add-flows 10
table=0,priority=777,in_port=121,actions=output:99,Set_field:ff:ff:ff:ff:55:66-\>eth_src

switch (config)# openflow add-flows 11 table=0,priority=777,in_port=127,actions=group:11

switch (config)# openflow add-flows 12 priority=12,in_port=105,actions=group:5

switch (config)# openflow add-flows 13 table=0,priority=777,in_port=127,actions=meter:6,output:117

switch (config)# openflow add-flows 14 table=2,priority=777,in_port=127,actions=meter:2,output:117

switch (config)# openflow add-flows 10 ip,priority=10,in_port=Eth1/1,dl_vlan=10,actions=output=Eth1/11

switch (config)# openflow add-flows 40
ip,priority=10,in_port=Eth1/1,action=set_field:00:0c:e9:00:00:01 eth_src,output=Eth1/11

switch (config)# openflow add-flows 30 ip,priority=100,actions=output=normal

switch (config)# openflow add-flows 10 priority=10,in_port=ANY,actions=DROP
```

Related Commands	

Notes	• If no flow-text is provided the command deletes the configured OpenFlow flows
	• The unique ID of the rule is the set of match keys, priority, and table number. Please notice that if match keys, priority, and table number are the same while the list of actions are different, it will be treated as the same rule.
	`openflow add-flows 12`
	`table=0,priority=123,in_port=ANY,actions=group:1,mod_dl_src:ff:ff:ff:ff:55:66,`
	`group:3`
	`openflow add-flows 13`
	`table=0,priority=123,in_port=ANY,actions=group:2,mod_dl_src:ff:ff:ff:ff:55:66,`
	`group:1`

12.14.6.4 openflow del-flows

	openflow del-flows [<flow-id>] Deletes OpenFlow flow.	
Syntax Description	flow-id	ID number to give this flow Range: 0-65535
Default	N/A	
Configuration Mode	config	
History	3.6.4006	
	3.9.1000	Updated note
Example	`switch (config)# openflow del-flows 1`	
Related Commands		
Notes	If flow ID "all" is provided, the command deletes all configured OpenFlow flows	

12.14.6.5 openflow add-group

	openflow add-group <group-id> <group-type> <bucket-parameter1>[,...,<bucket-parameterN>] Adds an OpenFlow group.	
Syntax Description	group-id	Group ID number
	group-type	For a list of supported group types, see the group column in "OpenFlow 1.3 Pipeline Capabilities Summary Table"
	bucket parameter	Possible values: • actions=output,...,output • bucket_id=<id-number> • watch_group=<group_id> • watch_port=<port> • weight=<value>
Default	N/A	
Configuration Mode	config	
History	3.6.4006	
	3.9.1600	Added note
Example	`switch (config) # openflow add-group group_id=3,type=ff,bucket=watch_port:117,output:123,bucket=watch_port:123,output:119,bucket=watch_port:111,output:119,113,121,115,123,109,117`	
Related Commands		
Notes	The maximum number of ports in one OpenFlow group is as follows: • 64 ports in Spectrum-based systems • 128 ports in Spectrum-2 and Spectrum-3 systems More than one group in the action list of OpenFlow is not supported	

12.14.6.6 openflow del-group

	openflow del-group <group-id> Deletes matching OpenFlow group ID.	
Syntax Description	group-id	Group ID number
Default	N/A	
Configuration Mode	config	
History	3.6.4006	
	3.9.1000	Updated note
Example	`switch (config)# openflow del-group`	
Related Commands		
Notes	If group ID "all" is provided, the command deletes all configured OpenFlow groups.	

12.14.6.7 openflow mod-group

	openflow mod-group <group-id> <group-type> <bucket-parameter1>[,...,<bucket-parameterN>] Modifies matching OpenFlow group ID.	
Syntax Description	group-id	Group ID number
	group-type	For a list of supported group types, see the group column in "OpenFlow 1.3 Pipeline Capabilities Summary Table"
	bucket parameter	Possible values: • actions=output,...,output • bucket_id=<id-number> • watch_group=<group_id> • watch_port=<port> • weight=<value>
Default	N/A	
Configuration Mode	config	
History	3.6.4006	
Example	`switch (config)# openflow mod-group group_id=3,type=ff,bucket=watch_port:117,output:123,bucket=watch_port:123,output:119,bucket=watch_port:111,output:119,113,121,115,123,109,117,119`	
Related Commands	openflow add-group	
Notes	A group must exist in order to execute this command	

12.14.6.8 openflow add-meter

	openflow add-meter <meter-id> <meter-rule> <band-parameter1>[,...,<band-parameterN>] Adds OpenFlow meter.	
Syntax Description	meter-id	Meter ID number

	meter-rule	For a list of supported meters types, see the meter column in "OpenFlow 1.3 Pipeline Capabilities Summary Table"
	band-parameter	Possible values: • type={type | drop} • rate=<value> • burst_size=<size>
Default	N/A	
Configuration Mode	config	
History	3.6.4006	
Example	`switch (config)# openflow add-meter meter=6,pktps,band=type=drop,rate=10`	
Related Commands		
Notes		

12.14.6.9 openflow del-meter

	openflow del-meter <meter-id> Deletes matching OpenFlow meter ID.	
Syntax Description	meter-id	Meter ID number
Default	N/A	
Configuration Mode	config	
History	3.6.4006	
	3.9.1000	Updated note
Example	`switch (config)# openflow del-meter meter=6`	
Related Commands		
Notes	If meter ID "all" is provided, the command deletes all configured OpenFlow meters.	

12.14.6.10 openflow fail-mode secure

	openflow fail-mode secure no openflow fail-mode secure Enables the "fail secure mode" of the switch. The no form of the command disables the "fail secure mode" of the switch.	
Syntax Description	N/A	
Default	Enabled	
Configuration Mode	Config	
History	3.8.2100	
	3.9.1600	Added note below
Example	`switch (config) # no openflow fail-mode secure`	
Related Commands		

Notes	In the case that a switch loses contact with all controllers as a result of echo request timeouts, TLS session timeouts, or other disconnections, the switch should immediately enter either “fail secure mode” or “fail standalone mode" (depending upon the switch implementation and configuration). "Fail secure mode" only affects the switch behavior in that packets and messages destined to go to the controllers are dropped. Flow entries should continue to expire according to their timeouts in “fail secure mode." In “fail standalone mode," the switch processes all packets using the OFPP_NORMAL reserved port and the switch acts as a legacy Ethernet switch or router.

Note that the default fail-mode is "secure". There is no default rule with action normal for this mode. All traffic will be affected, including protocols, until required rule is added or fail-mode is changed to "standalone". If using controller, add required rule via controller in any fail-mode.

12.14.6.11 openflow mod-meter

	openflow mod-meter <meter-id> <meter-rule> <band-parameter1>[,...,<band-parameterN>] Modifies matching OpenFlow meter ID.	
Syntax Description	meter-id	Meter ID number
	meter-rule	For a list of supported meters types, see the meter column in “OpenFlow 1.3 Pipeline Capabilities Summary Table”
	band-parameter	Possible values: • type={type | drop} • rate=<value> • burst_size=<size>
Default	N/A	
Configuration Mode	config	
History	3.6.4006	
Example	`switch (config)# openflow mod-meter meter=6,pktps,band=type=drop,rate=10`	
Related Commands		
Notes		

12.14.6.12 openflow re-apply flows

	openflow re-apply flows <flow-id> Reapplies matching flow ID.	
Syntax Description	flow-id	Range: 0-65535
Default	N/A	
Configuration Mode	config	
History	3.6.4006	
Example	`switch (config)# openflow re-apply flows 58`	
Related Commands		

Notes	

12.14.6.13 openflow re-apply groups

	openflow re-apply groups <group-id> Reapplies matching group ID.	
Syntax Description	group-id	Range: 0-65535
Default	N/A	
Configuration Mode	config	
History	3.6.4006	
Example	`switch (config)# openflow re-apply groups group_id=2`	
Related Commands		
Notes		

12.14.6.14 openflow re-apply meters

	openflow re-apply meters <meter-id> Reapplies matching meters ID.	
Syntax Description	meter-id	Range: 0-65535
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.4006	
Example	`switch (config interface ethernet 1/1)# openflow re-apply meters 13`	
Related Commands		
Notes		

12.14.6.15 controller-ip

	openflow controller-ip <ip-address> [tls] [tcp-port <tcp-port>] no openflow controller-ip <ip-address> [tls] [tcp-port <tcp-port>] Configures the OpenFlow controller's IP & TCP port. The command "no openflow controller-ip <ip-address>" deletes all OpenFlow controller configurations related to its IP address. The command "no openflow controller-ip <ip-address> tcp-port" deletes all the OpenFlow controller configurations related to IP address, and any tcp-port except for TLS ones. The command "no openflow controller-ip <ip-address> [tls] tcp-port <tcp-port>" deletes the entry for the OpenFlow controller IP address, TLS (if applicable), and the TCP port	
Syntax Description	ip-address	The IPv4 address of the OpenFlow controller
	tls	Configures secure connection to OpenFlow controller
	tcp-port	Sets the TCP port number of the OpenFlow controller
Default	TCP port 6633	

Configuration Mode	config openflow	
History	3.6.1002	
	3.6.2002	Added “tls” parameter
Example	`switch (config openflow) # controller-ip 10.10.10.10 tls tcp-port 6633`	
Related Commands		
Notes		

12.14.6.16 datapath-id

	datapath-id <value> no datapath-id Sets a specific identifier for the switch with which the controller is communicating. The no form of the command resets the parameter to its default value.	
Syntax Description	value	The most significant 16 bits of the agent data-path ID Range: 0x0000-0xFFFF in hexa
Default	0x0000	
Configuration Mode	config openflow	
History	3.3.4200	
Example	`switch (config openflow) # datapath-id 0x1234`	
Related Commands		
Notes		

12.14.6.17 openflow table match-keys

	openflow table <table_id[-table_id]> match-keys <key_list> no openflow table <table_id[-table_id]> match-keys [<key_list>] Adds ACL keys to an OpenFlow table. The no form of the command removes ACL keys from the OpenFlow table.	
Syntax Description	table_id	OpenFlow table ID for adding/removing key values. Can be one ID or range. Range: 0-249.
	key_list	Key value(s)
Default	0x0000	
Configuration Mode	config	
History	3.3.4200	
	3.9.0300	Added note of supported keys
Example	`switch (config) # openflow table 1 match-keys metadata ip_proto`	
Related Commands		

Notes	• OpenFlow match rules are installed according to the configured match keys • New match keys are configured only when the table is empty (i.e. does not contain any rules) • The following are the supported keys in this command: Key name Description in_port Source port dl_src Source MAC address dl_dst Destination MAC address dl_type Ethernet protocol type vlan_vid Virtual LAN tag vlan_pcp Priority Code Point ip_src Source IPv4 address ip_dst Destination IPv4 address ip_proto IPV4 - Next protocol, IPV6 - Next header ip_dscp IP ToS/DSCP or IPv6 traffic class field dscp ip_ecn ECN bits from IP header ip_ttl IP TTL or IPv6 hop limit l4_src_port Source L4 port l4_dst_port Destination L4 port

12.14.6.18 openflow acl table counter disable

	openflow acl table <id/range> counter disable no openflow acl table <id/range> counter disable Disables counter for a specific ACL or range of ACL OpenFlow tables. The no form of the command enables counter for ACL table.	
Syntax Description	id/range	Specific ACL or range of ACL OpenFlow tables.
Default	Enabled counter for all ACL tables.	
Configuration Mode	config	
History	3.9.2000	
Example	switch (config) # openflow acl table 10 counter disable switch (config) # openflow acl table 0-249 counter disable switch (config) # no openflow acl table 0-10 counter disable	
Related Commands		
Notes		

12.14.6.19 show openflow

	show openflow Displays general information about the OpenFlow protocol configuration.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4200	
	3.6.1002	Updated example
Example		

```
switch (config) # show openflow
OpenFlow Version: OpenFlow 1.3
Datapath ID: ffff7cfe90e600c0
Controllers Information:
Controller                State          Role       Changed (sec)  Last Error
-----------               -----          ----       -------------  ----------
tcp:1.1.1.1:6633          BACKOFF        other      3 Connection   timed out
tcp:10.10.10.10:6633      ACTIVE         other      2067 N/A
tcp:10.10.10.30:6633      ACTIVE         other      2067 N/A

Mapping of OpenFlow ports to their OpenFlow numbers:
Interface OF-Port
--------- -------
Eth1/12   OF107
Eth1/9    OF109
Eth1/10   OF111
Eth1/7    OF113
Eth1/8    OF115
Eth1/3    OF121
Eth1/4    OF123
```

Related Commands	
Notes	

12.14.6.20 show openflow flows

	show openflow flows Displays information about the OpenFlow flows.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4302	
	3.6.1002	Updated example

Example	`switch (config) # show openflow flows`
	`OFPST_FLOW reply (OF1.3) (xid=0x2):`
	`cookie=0x0, duration=467.993s, table=0, n_packets=0, n_bytes=0, send_flow_rem priority=8,in_port=125 actions=output:123`
	`cookie=0x0, duration=439.218s, table=0, n_packets=0, n_bytes=0, send_flow_rem priority=9999,in_port=125 actions=output:123`
	`cookie=0x0, duration=467.984s, table=0, n_packets=0, n_bytes=0, send_flow_rem priority=1000 actions=drop`
	`cookie=0x0, duration=467.975s, table=0, n_packets=0, n_bytes=0, send_flow_rem priority=200,dl_vlan=222 actions=pop_vlan,output:123`
	`cookie=0x0, duration=467.987s, table=0, n_packets=0, n_bytes=0, send_flow_rem priority=10,dl_vlan=10 actions=output:123`
	`cookie=0x0, duration=468.013s, table=0, n_packets=0, n_bytes=0, send_flow_rem priority=8,dl_dst=01:01:01:01:01:01 actions=output:123`
	`cookie=0x0, duration=467.991s, table=0, n_packets=0, n_bytes=0, send_flow_rem priority=8,dl_src=01:01:01:01:01:01 actions=output:123`
	`cookie=0x0, duration=467.992s, table=0, n_packets=0, n_bytes=0, send_flow_rem priority=5,arp actions=output:123`	

Related Commands		
Notes		

12.14.6.21 show openflow flows ethernet-names

	show openflow flows <cookie | table> ethernet-names Displays OpenFlow flows configuration with interface names.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4006	
	3.9.2400	Updated example with LAG and MLAG interfaces
Example		

```
switch (config) # show openflow flows ethernet-names
OFPST_FLOW reply (OF1.3) (xid=0x2):
cookie=0x0, duration=911.531s, table=0, n_packets=0, n_bytes=0, priority=0 actions=NORMAL
cookie=0x0, duration=80.662s, table=1, n_packets=0, n_bytes=0,
priority=0,in_port=0,dl_src=02:00:00:00:00:00 actions=output:Eth1/13,output:123,output:127
cookie=0x0, duration=80.530s, table=1, n_packets=0, n_bytes=0,
priority=1,in_port=1,dl_src=02:01:00:00:00:00 actions=output:Eth1/13,output:123,output:127
cookie=0x0, duration=80.414s, table=1, n_packets=0, n_bytes=0,
priority=2,in_port=2,dl_src=02:02:00:00:00:00 actions=output:Eth1/13,output:123,output:127
cookie=0x0, duration=80.296s, table=1, n_packets=0, n_bytes=0,
priority=3,in_port=3,dl_src=02:03:00:00:00:00 actions=output:Eth1/13,output:123,output:127
cookie=0x0, duration=80.180s, table=1, n_packets=0, n_bytes=0,
priority=4,in_port=4,dl_src=02:04:00:00:00:00 actions=output:Eth1/13,output:123,output:127
cookie=0x0, duration=80.064s, table=1, n_packets=0, n_bytes=0,
priority=5,in_port=5,dl_src=02:05:00:00:00:00 actions=output:Eth1/13,output:123,output:127
cookie=0x0, duration=79.948s, table=1, n_packets=0, n_bytes=0,
priority=6,in_port=6,dl_src=02:06:00:00:00:00 actions=output:Eth1/13,output:123,output:127
cookie=0x0, duration=79.831s, table=1, n_packets=0, n_bytes=0,
priority=7,in_port=7,dl_src=02:07:00:00:00:00 actions=output:Eth1/13,output:123,output:127
cookie=0x0, duration=79.711s, table=1, n_packets=0, n_bytes=0,
priority=8,in_port=8,dl_src=02:08:00:00:00:00 actions=output:Eth1/13,output:123,output:127
cookie=0x0, duration=79.591s, table=1, n_packets=0, n_bytes=0,
priority=9,in_port=9,dl_src=02:09:00:00:00:00 actions=output:Eth1/13,output:123,output:127
cookie=0x0, duration=79.467s, table=1, n_packets=0, n_bytes=0,
priority=10,in_port=10,dl_src=02:0a:00:00:00:00 actions=output:Eth1/13,output:123,output:127
cookie=0x0, duration=79.445s, table=1, n_packets=0, n_bytes=0,
priority=10,in_port=Mpo1,actions=output:Eth1/13,Po2
cookie=0x0, duration=79.445s, table=1, n_packets=0, n_bytes=0,
priority=10,in_port=Eth1/12,actions=output:Mpo2
```

Related Commands	
Notes	

12.14.6.22 show openflow groups

	show openflow groups Displays OpenFlow flows configuration with interface names.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.3004
Example	

```
switch (config) # show openflow groups
OFPST_GROUP_DESC reply (OF1.3) (xid=0x2):
group_id=5566,type=select,bucket=weight:5,actions=output:1,bucket=weight:7,actions=output:2,bucket=weight:
22,actions=output:3
```

Related Commands	
Notes	

12.14.6.23 show openflow groups ethernet-names

	show openflow groups ethernet-names Displays all the configured OpenFlow groups with their interface names.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4006	
	3.9.2400	Updated Example with LAG and MLAG interfaces
Example		

```
switch (config) # show openflow groups
OFPST_GROUP_DESC reply (OF1.3) (xid=0x2):
group_id=4,type=all,bucket=actions=output:Eth1/13,output:123
group_id=1,type=select,bucket=actions=output:Eth1/7,output:Eth1/8,output:Eth1/5,output:123,set_field:11:22
:33:44:00:00->eth_dst
group_id=2,type=select,bucket=actions=output:Eth1/13
group_id=3,type=all,bucket=actions=output:Eth1/13,output:123,set_field:11:22:33:44:00:00->eth_dst
group_id=5,type=all,bucket=actions=output:Mpo1,output:Eth1/12
group_id=6,type=all,bucket=actions=output:Eth1/7,Mpo1,output:Eth1/5,set_field:
11:22:33:44:00:00->eth_dst
```

Related Commands	
Notes	

12.14.6.24 show openflow meters

	show openflow meters [<ID>] Displays all/specified OpenFlow meters.	
Syntax Description	ID	Requested meter ID
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Example		

```
switch (config) # show openflow meters
OFPST_METER_CONFIG reply (OF1.3) (xid=0x2):
meter=20 kbps bands=
type=drop rate=300

meter=100 kbps bands=
type=drop rate=500

meter=200 kbps bands=
type=drop rate=500

switch (config) # show openflow meters 20
OFPST_METER_CONFIG reply (OF1.3) (xid=0x2):
meter=20 kbps bands=
type=drop rate=300
```

Related Commands	
Notes	

12.14.6.25 show openflow flows table

	show openflow flows table <NUM> [summary] Displays information/summary of a given OpenFlow flows table.	
Syntax Description	NUM	NUM range: 0-252
	summary	Displays given OpenFlow flow table summary
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Example		

```
switch (config) # show openflow flows table 1
OFPST_FLOW reply (OF1.3) (xid=0x2):
cookie=0x0, duration=6.344s, table=1, n_packets=0, n_bytes=0, in_port=127 actions=drop

switch (config) # show openflow flows table 1 summary
OFPST_AGGREGATE reply (OF1.3) (xid=0x2): packet_count=0 byte_count=0 flow_count=1
```

Related Commands	
Notes	

12.14.6.26 show openflow flows cookie

	show openflow flows cookie <cookie> [summary] Displays information/summary of a given OpenFlow flows cookie.	
Syntax Description	cookie	Requested cookie ID in the following format: cookie_id.cookie_id/mask_id (e.g., 0x2A, 0x12/0x2)
	summary	Displays given OpenFlow flow table summary
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Example		

```
switch (config) # show openflow flows cookie 0x11
OFPST_FLOW reply (OF1.3) (xid=0x2):
cookie=0x11, duration=2.699s, table=0, n_packets=0, n_bytes=0, actions=NORMAL
switch (config) # show openflow flows cookie 0x22
OFPST_FLOW reply (OF1.3) (xid=0x2):
cookie=0x22, duration=3.970s, table=1, n_packets=0, n_bytes=0, in_port=127 actions=drop
```

Related Commands	
Notes	A cookie may be associated with a flow using the add-flows, and mod-flows commands.

12.14.6.27 show openflow table match-keys

	show openflow table <table_id[-table_id]> match-keys Displays configured ACL keys in OpenFlow table.

Syntax Description	table_id	OpenFlow table ID for adding/removing key values. Can be one ID or range. Range: 0-249.
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Example	switch (config) # show openflow table 2 match-keys Table: 2 Pending keys: Key name Description --------- ------------ in_port Source port dl_src Source MAC address dl_dst Destination MAC address dl_type Ethernet protocol type vlan_vid Virtual LAN tag vlan_pcp Priority Code Point ip_src Source IPv4 address ip_dst Destination IPv4 address ip_proto IPV4 - Next protocol, IPV6 - Next header ip_dscp IP ToS/DSCP or IPv6 traffic class field dscp ip_ecn ECN bits from IP header ip_ttl IP TTL or IPv6 hop limit l4_src_port Source L4 port l4_dst_port Destination L4 port metadata Matches value in the metadata field	
Related Commands		
Notes		

12.14.6.28 show openflow table match-keys supported

	show openflow table <table_id[-table_id]> match-keys supported Displays list of ACL keys which can be configured in OpenFlow table.	
Syntax Description	table_id	OpenFlow table ID for adding/removing key values. Can be one ID or range. Range: 0-249.
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Example	switch (config) # show openflow table 2 match-keys supported Key name Description --------- ------------ in_port Source port dl_src Source MAC address dl_dst Destination MAC address dl_type Ethernet protocol type vlan_vid Virtual LAN tag vlan_pcp Priority Code Point ip_src Source IPv4 address ip_dst Destination IPv4 address ipv6_dst Destination IPv6 address ipv6_src Source IPv6 address ip_proto IPV4 - Next protocol, IPV6 - Next header ip_dscp IP ToS/DSCP or IPv6 traffic class field dscp ip_ecn ECN bits from IP header ip_ttl IP TTL or IPv6 hop limit l4_src_port Source L4 port l4_dst_port Destination L4 port metadata Matches value in the metadata field	

Related Commands	
Notes	

13 VXLAN

Data centers are being increasingly consolidated and outsourced in an effort to improve the deployment time of applications and reduce operational costs, and applications are constantly raising demand for compute, storage, and network resource. Thus, in order to scale compute, storage, and network resources, physical resources are being abstracted from their logical representation, in what is referred to as server, storage, and network virtualization. Virtualization can be implemented in various layers of computer systems or networks.

Multi-tenant data centers are taking advantage of the benefits of server virtualization to provide a new kind of hosting—a virtual hosted data center. Multi-tenant data centers are ones where individual tenants could belong to a different company or a different department. To a tenant, virtual data centers are similar to their physical counterparts, consisting of end-stations attached to a network, complete with services such as load balancers and firewalls. To tenant systems, a virtual network looks like a normal network, except that the only end-stations connected to the virtual network are those belonging to a tenant's specific virtual network.

How a virtual network is implemented does not generally matter to the tenant; what matters is that the service provided (Layer 2 (L2) or Layer 3 (L3)) has the right semantics, performance, etc. It could be implemented via a pure routed network, a pure bridged network, or a combination of bridged and routed networks.

VXLAN (Virtual eXtensible Local Area Network) addresses the above requirements of the L2 and L3 data center network infrastructure in the presence of virtual networks in a multi-tenant environment. It runs over the existing networking infrastructure and provides a means to "stretch" an L2 network. Each overlay bridge is called a VXLAN segment. Only machines within the same VXLAN segment can communicate with each other. Each VXLAN segment is identified through a 24-bit segment ID called "VXLAN Network Identifier (VNI)". A network endpoint which performs a conversion from virtual to physical network and back is called VXLAN Tunnel End-Point or VTEP.

In virtual environments, it is typically required to use logical switches to forward traffic between different virtual machines (VMs) on the same physical host, between virtual machines and the physical machines and between networks. Virtual switch environments use an OVSDB management protocol for configuration and state discovery of the virtual networks. OVSDB protocol allows programmable access to the database of virtual switch configuration.

13.1 Configuring VXLAN

To enable VXLAN:

1. Configure jumbo frames for NVE ports. Run:

```
switch (config)# interface ethernet 1/1-1/4 mtu 9216 force
```

2. Configure jumbo frames for underlay-facing ports. Run:

```
switch (config)# interface ethernet 1/17 mtu 9216 force
```

3. Create VLAN for all VXLAN traffic. Run:

```
switch (config)# vlan 3
```

4. Configure Overlay interfaces with VXLAN VLAN. Run:

```
switch (config)# interface ethernet 1/17 switchport access vlan 3
```

5. Enable IP routing. Run:

```
switch (config)# ip routing vrf default
```

6. Configure interface on the VXLAN VLAN and configure an IP address for it. Run:

```
switch (config)# interface vlan 3
switch (config interface vlan 3)# ip address 33.33.33.254 255.255.255.0
switch (config interface vlan 3)# interface vlan 3 mtu 9216
```

7. Enable NVE protocol. Run:

```
switch (config)# protocol nve
```

8. Configure interface NVE. Run:

```
switch (config)# interface nve 1
```

9. Create loopback interface to terminate the VXLAN tunnel. The IP address of the interface will be a VTEP endpoint address, and needs to be reachable in the underlay network. Run:

```
switch (config)# interface loopback 1
switch (config interface loopback 1)# ip address 1.2.3.4 255.255.255.255
switch (config)# interface nve 1 vxlan source interface loopback 1
```

10. Configure routing to other VTEP devices. Run:

```
switch (config)# ip route vrf default 1.2.3.5 /32 33.33.33.253
switch (config)# ip route vrf default 1.2.3.6 /32 33.33.33.252
```

11. Configure overlay-facing ports for NVE mode. Run:

```
switch (config)# interface ethernet 1/1 nve mode only force
switch (config)# interface ethernet 1/2 nve mode only force
switch (config)# interface ethernet 1/3 nve mode only force
switch (config)# interface ethernet 1/4 nve mode only force
```

For deployments with a controller, set up OVSDB:

1. Start OVSDB server. Run:

```
switch (config)# ovs ovsdb server
```

2. Configure the OVSDB manager to an IP address of a controller. Run:

```
switch (config)# ovs ovsdb manager remote ssl ip address 10.130.250.5
```

For controller-less deployments, configure the bridging from the CLI directly:

1. Create bridges. Run:

```
switch (config)# interface nve 1 nve bridge 7777
switch (config)# interface ethernet 1/1 nve vlan 10 bridge 7777
```

2. Configure source-node replication. Run:

```
switch (config)# no interface nve 1 nve fdb flood load-balance
```

3. Configure flood addresses for BUM traffic. Run:

```
switch (config)# interface nve 1 nve fdb flood bridge 7777 address 1.2.3.5
switch (config)# interface nve 1 nve fdb flood bridge 7777 address 1.2.3.6
```

4. Configure FDB remote learning. Run:

```
switch (config)# interface nve 1 nve fdb learning remote
```

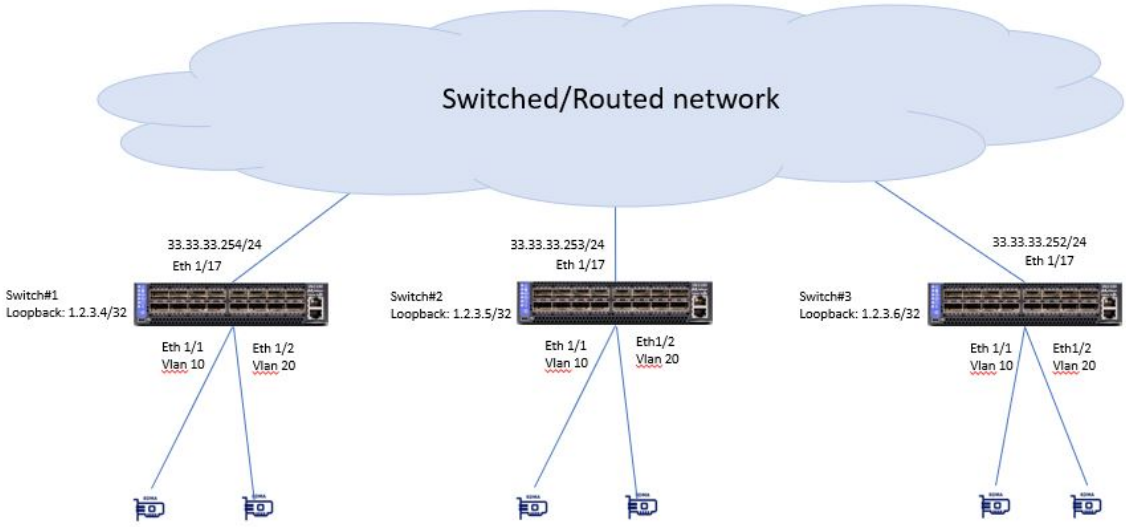

13.2 VMware Network Virtualization and Security Platform (NSX) Configuration

13.2.1 Hardware Topology

- 2 ESXi servers pre-configured with VXLAN networking using VMware NSX
- 3 NSX Controllers available for VXLAN unicast type logical switches
- 1 NVIDIA switch connected to the ESXi servers and to a physical database server
- Out-of-band network for management and a VLAN network to carry VXLAN traffic

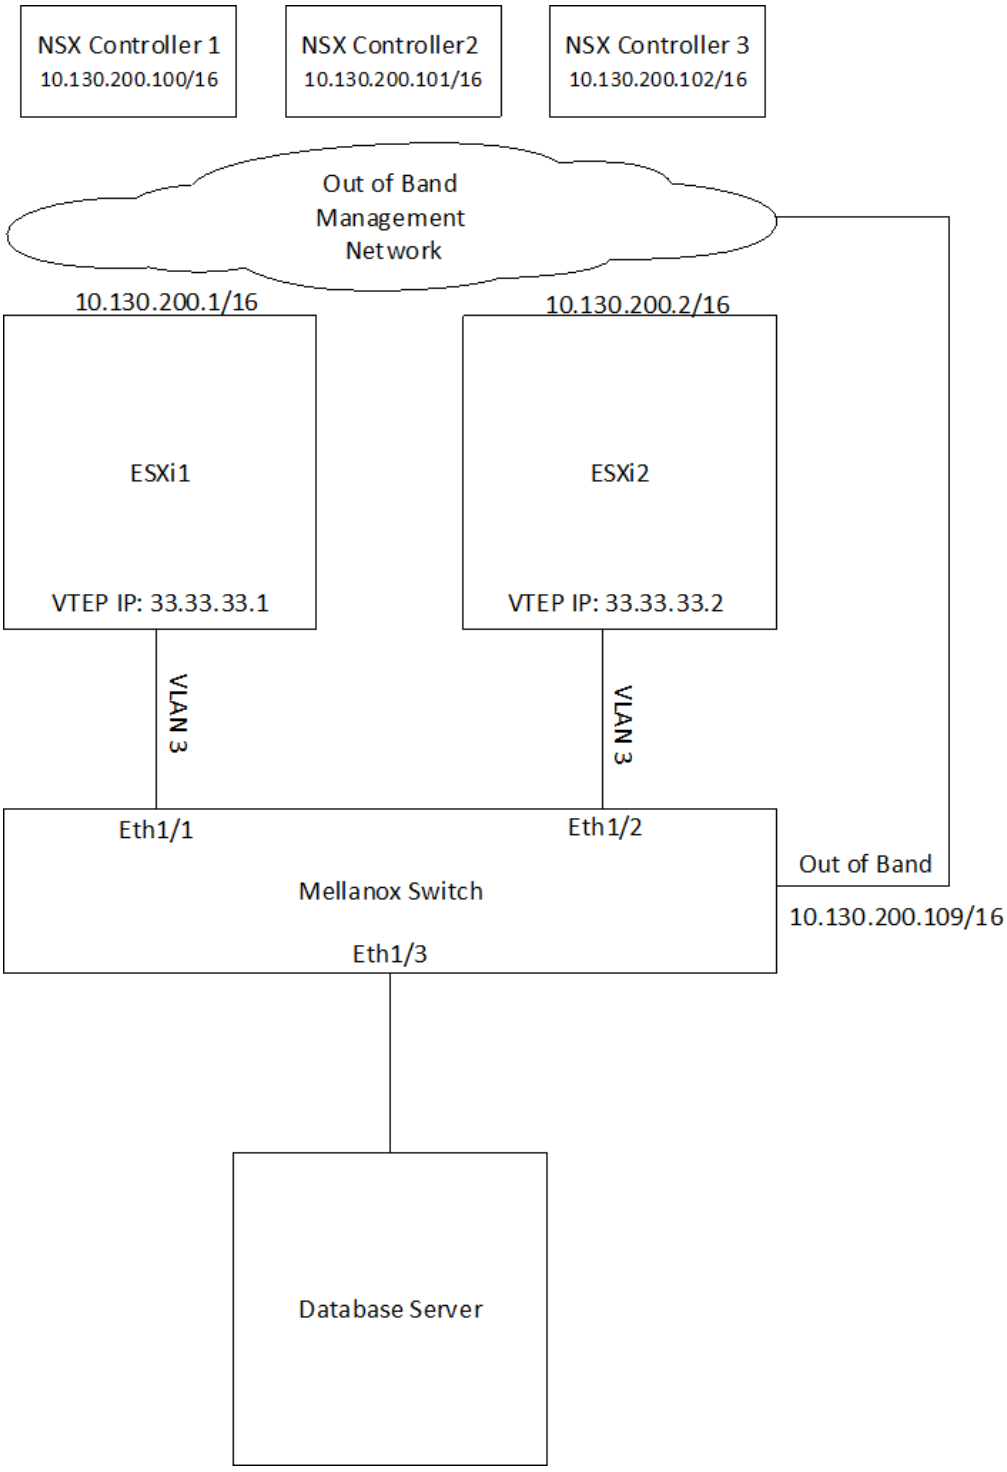

13.2.2 Switch Configuration

1. Configure jumbo frames on ESXi and Database server facing interfaces. Run:

```
switch (config)# interface ethernet 1/1-1/3 mtu 9216 force
```

2. Create VLAN 3 to carry VXLAN traffic (if it does not exist yet). Run:

```
switch (config)# vlan 3
switch (config vlan 3)# exit
switch (config)#
```

3. Enable IP routing. Run:

```
switch (config)# ip routing vrf default
```

4. Create an interface on VLAN 3 and assign an IP address to it.
   The IP address must be the default gateway of the VXLAN netstack created by NSX after enabling VXLAN traffic on the hosts.
   To check the default gateway in vSphere web client select an ESXi host and go to: Configure -> TCP/IP configuration.

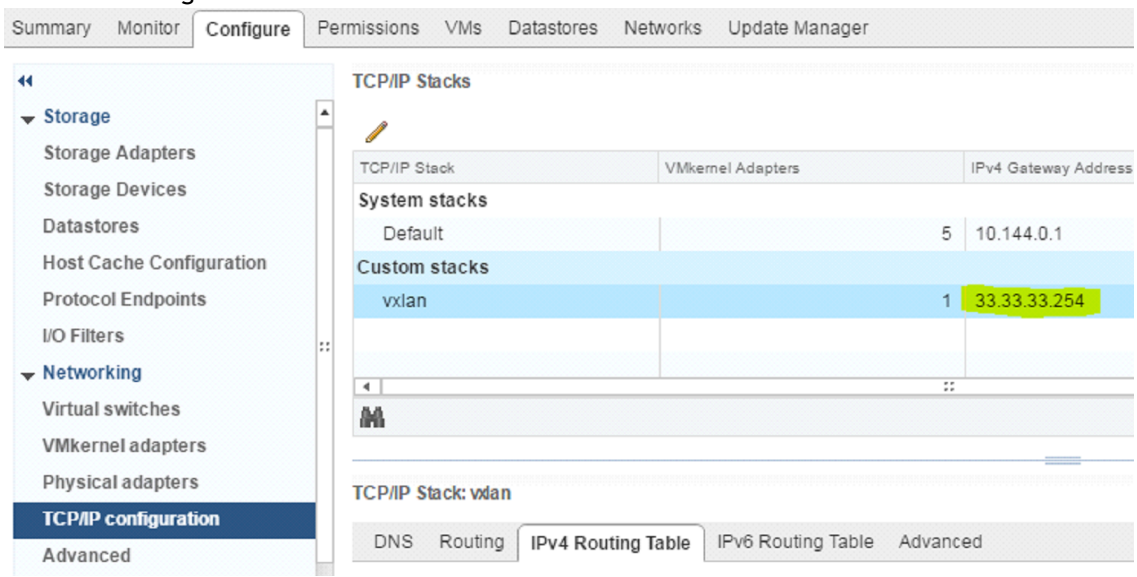


```
switch (config)# interface vlan 3
switch (config interface vlan 3)# ip address 33.33.33.254 255.255.255.0
switch (config interface vlan 3)# interface vlan 3 mtu 9216
```

5. Create a loopback interface to communicate with VTEPs on the ESXi servers by routing through “interface vlan 3”. This interface will be the VTEP IP assigned to the switch. Run:

```
switch (config)# interface loopback 1
switch (config interface loopback 1)# ip address 1.2.3.4 255.255.255.255
```

6. Enable NVE protocol. Run:

```
switch (config)# protocol nve
```

7. Configure interface NVE. Run:

```
switch (config)# interface nve 1
```

8. Configure the source of the NVE interface to be the loopback created above. Run:

```
switch (config)# interface nve 1 vxlan source interface loopback 1
```

9. Start the OVSDB server and connect it to the NSX Controllers. Run:

```
switch (config)# ovs ovsdb server
switch (config)# ovs ovsdb manager remote ssl ip address 10.130.200.100
switch (config)# ovs ovsdb manager remote ssl ip address 10.144.200.101
switch (config)# ovs ovsdb manager remote ssl ip address 10.144.200.102
```

10. Configure the port facing the Database server as an NVE port. Run:

```
switch (config)# interface ethernet 1/3 nve mode only force
```

11. Get the switch certificate for later configuration in the NSX Manager. Run:

```
switch (config)# show crypto certificate name system-self-signed public-pem
```

Copy the certificate starting with the line:

```
-----BEGIN CERTIFICATE-----
```

Until the line:

```
-----END CERTIFICATE-----
```

Make sure to include both of those lines.

NSX Manager Configuration

Adding Hosts to Replication Cluster

12. In NSX Manager, go to “Service Definitions” → “Hardware Devices”.

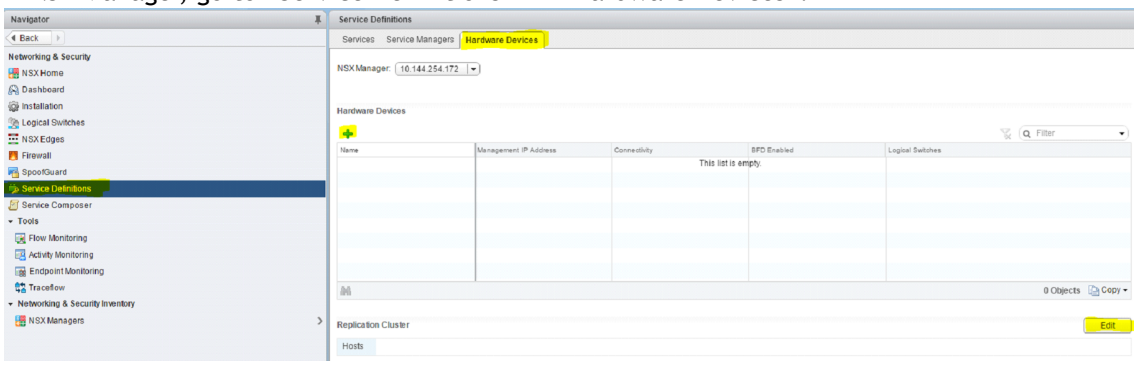

13. Under “Replication Cluster” click Edit.
14. Add both of the ESXi servers to the replication cluster.

All hosts added to the replication cluster can replicate BUM (Broadcast, Unknown unicast and Multicast) traffic to other ESXi servers.

When the switch needs to send BUM traffic to a virtual machine, it will select one of the hosts in the replication cluster and send the traffic to it, the host will then replicate it to all other ESXi hosts.

It is recommended to add at least 2 ESXi servers to the replication cluster for redundancy.

13.2.3 Adding the Switch to NSX

1. Under Hardware Devices click the + sign to add a new hardware device.

2. Fill in a name for the new hardware device.
3. Fill in the switch certificate we got earlier.
4. Click OK.

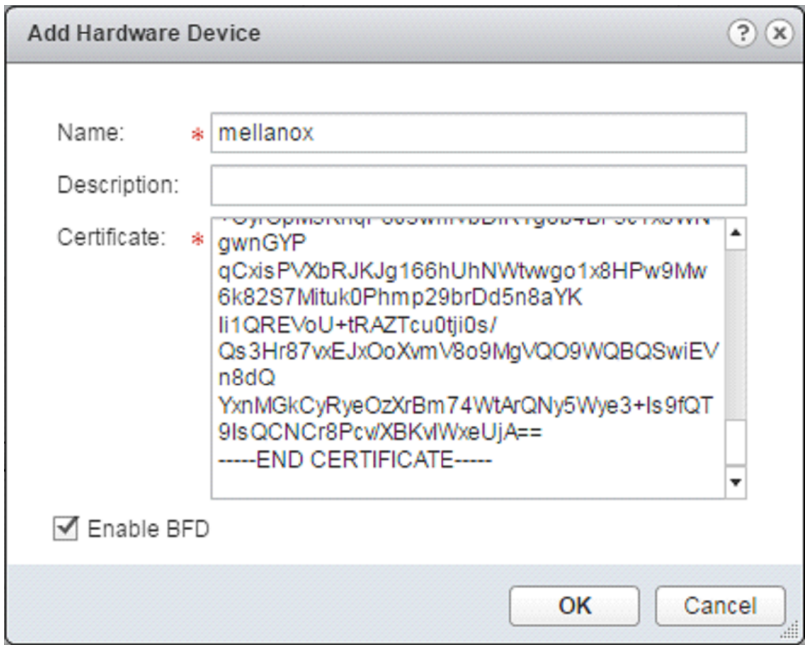

5. Wait until the new switch is showing as “UP” under the connectivity column, you may need to refresh vSphere client a few times.

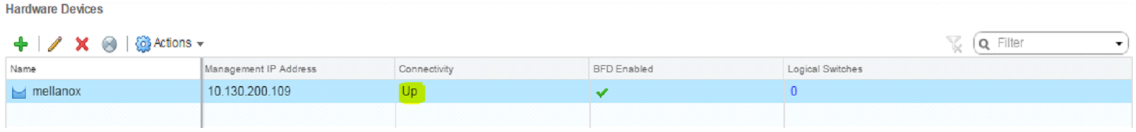

13.2.4 Mapping a Logical Switch to a Physical Switch Port

1. In NSX Manager go to “Logical Switches”.
2. Right click the logical switch you wish to map to the physical switch port and select “Manage Hardware Bindings”.

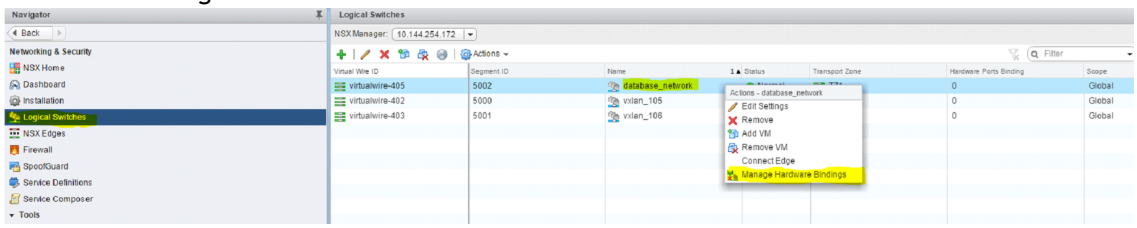

3. Click the “+” sign to add a new mapping instance.
4. Click Select under the port column and select port “eth3”, this corresponds to “ 1/3” we configured earlier as an NVE port in the switch.
5. Under the VLAN column, set the VLAN that will map this logical switch to this specific switch port, you can have multiple logical switches mapped to the same port on a different VLAN (for example to connect a firewall appliance to logical switches). For “access” configuration (no VLAN is required on the host connected to the physical switch port) use VLAN 1.

6. Click OK.

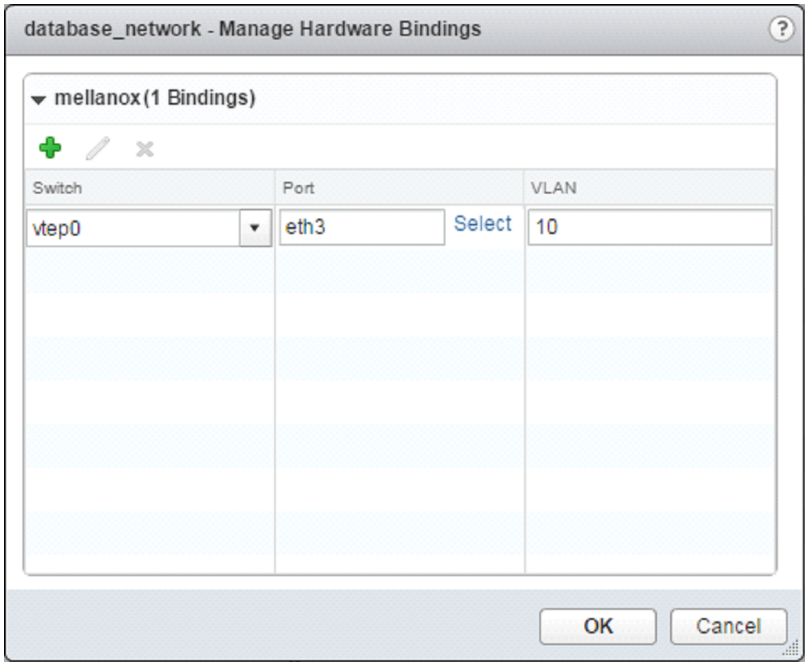

13.3 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community posts:

- HowTo Configure Openstack L2 Gateway with Spectrum Switch (VTEP)
- HowTo Configure VTEP using VMware NSX on Spectrum Switches

13.4 RoCE Over VXLAN

13.4.1 RoCEv2 Using PFC and ECN

The following figure and flow demonstrate how to configure RoCEv2 using PFC and ECN. RoCEv2 QoS is preserved by DSCP.

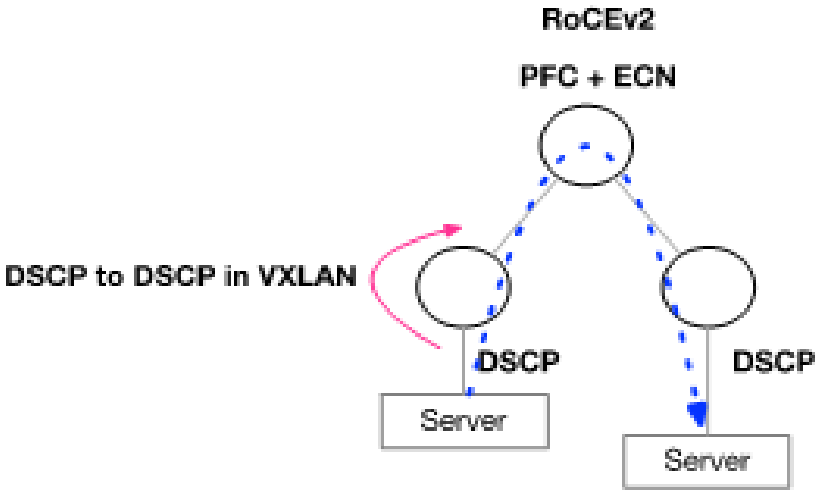

DSCP is automatically driven from the original packet into the VXLAN header in Onyx.

- Configure the switch buffer to support lossless traffic.

```
traffic pool roce type lossless
traffic pool roce memory percent 50.00
traffic pool roce map switch-priority 3
```

- Enable ECN.

```
interface ethernet 1/15 traffic-class 3 congestion-control ecn minimum-absolute 150 maximum-absolute 1500
interface ethernet 1/16 traffic-class 3 congestion-control ecn minimum-absolute 150 maximum-absolute 1500
interface mlag-port-channel 7-8 traffic-class 3 congestion-control ecn minimum-absolute 150 maximum-
absolute 1500
interface port-channel 1 traffic-class 3 congestion-control ecn minimum-absolute 150 maximum-absolute 1500
interface ethernet 1/15 traffic-class 6 dcb ets strict
interface ethernet 1/16 traffic-class 6 dcb ets strict
interface mlag-port-channel 7-8 traffic-class 6 dcb ets strict
interface port-channel 1 traffic-class 6 dcb ets strict
```

- Set QoS trust to DSCP.

```
interface ethernet 1/15-1/16 qos trust L3
interface mlag-port-channel 7-8 qos trust L3
interface port-channel 1 qos trust L3
```

13.4.2 RoCEv1 Using PFC

The following figure and flow demonstrate how to configure RoCEv1 using PFC. RoCEv1 QoS is based on the PCP field sent by the server.

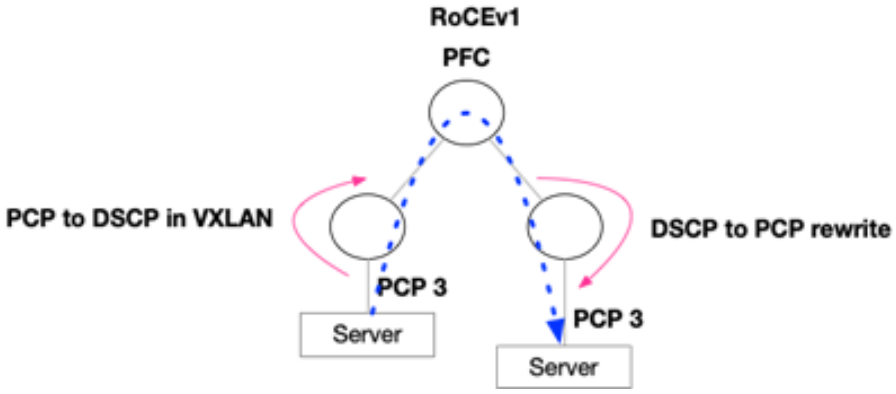

- Configure the switch buffer to support lossless traffic.

```
traffic pool roce type lossless
traffic pool roce memory percent 50.00
traffic pool roce map switch-priority 3
```

- Set Uplinks and IPL trust to DSCP.

```
interface ethernet 1/15-1/16 qos trust L3
interface port-channel 1 qos trust L3
```

- Set Downlinks trust to PCP.

```
interface mlag-port-channel 7-8 qos trust L2
```

- Set Downlinks rewrite to DSCP. This will allow translation from PCP to DSCP in VXLAN.

```
interface mlag-port-channel 7-8 qos rewrite dscp
```

- Set Uplinks and IPL rewrite to PCP. This will allow translation from DSCP to PCP.

```
interface ethernet 1/15-1/16 qos rewrite pcp
interface port-channel 1 qos rewrite pcp
```

13.5 VXLAN Commands

- VXLAN Commands

13.6 VXLAN Commands

- 13.6.1 protocol nve
- 13.6.2 interface nve
- 13.6.3 nve bridge
- 13.6.4 nve controller bgp
- 13.6.5 nve fdb flood bridge address

- 13.6.6 nve fdb flood load-balance
- 13.6.7 nve fdb learning remote
- 13.6.8 nve mode only
- 13.6.9 nve neigh-suppression
- 13.6.10 nve vlan bridge
- 13.6.11 nve vlan neigh-suppression
- 13.6.12 nve vni vlan
- 13.6.13 interface nve auto-vlan-map
- 13.6.14 interface nve disable nve vni
- 13.6.15 vxlan mlag-tunnel-ip
- 13.6.16 vxlan source interface loopback
- 13.6.17 shutdown
- 13.6.18 clear mac-address-table nve
- 13.6.19 clear nve counters
- 13.6.20 show interfaces nve
- 13.6.21 show interfaces nve detail
- 13.6.22 show interfaces nve counters
- 13.6.23 show interfaces counters vlan
- 13.6.24 show interfaces nve flood
- 13.6.25 show interfaces nve mac-address-table
- 13.6.26 show interfaces nve mac-address-table local learned unicast
- 13.6.27 show interfaces nve mac-address-table remote configured multicast
- 13.6.28 show interfaces nve peers
- 13.6.29 ovs ovsdb server
- 13.6.30 ovs ovsdb manager remote
- 13.6.31 ovs ovsdb server listen
- 13.6.32 ovs logging level
- 13.6.33 show ovs

13.6.1 protocol nve

	protocol nve no protocol nve Enables NVE functionality and displays NVE commands. The no form of the command hides the NVE commands and deletes its database.
Syntax Description	N/A
Default	no protocol nve
Configuration Mode	config
History	3.6.3004
Example	`switch (config) # protocol nve`
Related Commands	
Notes	

13.6.2 interface nve

	interface nve <nve-id> no interface nve <nve-id> Creates VXLAN tunnel. The no form of the command destroys VXLAN tunnel.	
Syntax Description	nve-id	NVE ID Range: 1-64
Default	N/A	
Configuration Mode	config	
History	3.6.3004	
Example	`switch (config) # interface nve 1` `switch (config interface nve 1) #`	
Related Commands	protocol nve	
Notes		

13.6.3 nve bridge

	nve bridge <vni-id> [name <bridge-name>] no nve bridge <vni-id> Creates an NVE bridge with a given VNI. The no form of the command removes NVE bridge.	
Syntax Description	vni-id	VXLAN network identifier Range: 0-16777216
	bridge-name	Name of NVE bridge
Default	bridge-name: bridge-<vni-id>	
Configuration Mode	config interface nve	
History	3.6.3212	
Example	`switch (config interface nve 1) # nve bridge 25`	
Related Commands	protocol nve	
Notes	Number of bridges limited to 500	

13.6.4 nve controller bgp

	nve controller bgp no nve controller bgp Enables the NVE controller mode to BGP. The no form disables the NVE controller mode from BGP to OVSDB mode.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface nve
History	3.8.1000

Example	`switch (config interface nve 1) # nve controller mode`
Related Commands	protocol nve
Notes	If controller BGP is enabled, shutdown command is not supported.

13.6.5 nve fdb flood bridge address

	nve fdb flood bridge <vni-id> address <ip-address> no nve fdb flood bridge <vni-id> address [ip-address] Adds an IP address of a remote VTEP to be used for BUM traffic. The no form of the command has two input options: • Entering an IP address removes a specific remote address • No IP address removes all addresses	
Syntax Description	vni-id	VXLAN network identifier Range: 0-16777216
	ip-address	IP address
Default	N/A	
Configuration Mode	config interface nve	
History	3.6.3212	
Example	`switch (config interface nve 1) # nve fdb flood bridge 7777 address 1.2.3.6`	
Related Commands	protocol nve	
Notes	The number of IP addresses is limited to 750	

13.6.6 nve fdb flood load-balance

	nve fdb flood load-balance no nve fdb flood load-balance Configures service-node replication. The no form of the command configures source-node replication.
Syntax Description	N/A
Default	service-node replication
Configuration Mode	config interface nve
History	3.6.8008
Example	`switch (config interface nve 1) # nve fdb flood load-balance`
Related Commands	protocol nve
Notes	

13.6.7 nve fdb learning remote

	nve fdb learning remote no nve fdb learning remote Enables remote (controller-less) FDB learning. The no form of the command disables remote FDB learning.

Syntax Description	N/A
Default	Disabled (controller-based learning)
Configuration Mode	config interface nve
History	3.6.8008
Example	`switch (config interface nve 1) # nve fdb learning remote`
Related Commands	protocol nve
Notes	

13.6.8 nve mode only

	nve mode only [force] no nve mode only [force] Sets physical interface to NVE mode. The no form of the command removes physical interface from NVE mode.	
Syntax Description	force	Forces configuration while interface is admin up
Default	no nve mode only	
Configuration Mode	config interface ethernet	
History	3.6.3004	
Example	`switch (config interface ethernet 1/1) # nve mode only`	
Related Commands	protocol nve	
Notes		

13.6.9 nve neigh-suppression

	nve neigh-suppression no nve neigh-suppression Enables neighbor suppression for all VLAN-VNI mappings. The no form of the command disables neighbor suppression for all VLAN-VNI mappings.	
Syntax Description	N/A	
Default	no nve mode only	
Configuration Mode	config interface nve	
History	3.8.1000	
	3.9.1000	Added support for IPv6 neighbor suppression
Example	`switch (config interface nve 1) # nve neigh-suppression`	
Related Commands	protocol nve nve controller bgp nve vlan neigh-suppression	
Notes	• If VLAN mapping is already configured, then the user might run "disable nve vlan <vlan_id> neigh-suppression" to not use global configuration. • BGP controller mode must be set prior to using this command	

13.6.10 nve vlan bridge

	nve vlan <vlan-id> bridge <vni-id> no nve vlan <vlan-id> bridge <vni-id> Maps a VLAN to a specific bridge on the interface (controller-less configuration). The no form of the command unmaps a VLAN from a specific bridge on the interface.	
Syntax Description	vni-id	VXLAN network identifier Range: 0-16777216
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.6102	
Example	`switch (config interface ethernet 1/1) # nve vlan 10 bridge 7777`	
Related Commands	protocol nve	
Notes	• Multiple VLANs cannot be mapped to a single bridge • If you use VTEP light, VLAN 0 should be used for untagged traffic	

13.6.11 nve vlan neigh-suppression

	nve vlan <vlan_id> neigh-suppression [disable | no] nve vlan <vlan_id> neigh-suppression Configures neigh-suppression for a specific VLAN mapping. The no form of the command uses the global neigh-suppression configuration in this VLAN mapping. The disable form of the command disables neigh-suppression in this VLAN mapping regardless of the global configuration.	
Syntax Description	vlan_id	VXLAN network identifier Range: 1-4094
Default	N/A	
Configuration Mode	config interface nve	
History	3.8.1000	
Example	`switch (config interface nve 1) # nve vlan 5 neigh-suppression`	
Related Commands	protocol nve nve controller bgp nve neigh-suppression	
Notes	• BGP controller mode must be set prior to using this command • VLAN-VNI mapping needs to be set prior to running this command	

13.6.12 nve vni vlan

	nve vni <vni_value> vlan <vlan_id> [counter <encap/decap/both>] no nve vni <vni_value> vlan <vlan_id> Creates new VNI-to-VLAN manual mapping. The no form of the command deletes VNI-to-VLAN manual mapping.	
Syntax Description	vni_value	Possible values: 1-16777214

	vlan_id	VLAN ID Range: 1-4094
	encap	Enable counters for encapsulated packets per VLAN
	decap	Enable counters for decapsulated packets per VLAN
Default	N/A	
Configuration Mode	config interface nve	
History	3.8.1000	
	3.9.1000	Updated example and added counters per VLAN
Example	`switch (config interface nve 1) # nve vni 5000 vlan 5`	
Related Commands	protocol nve nve controller bgp interface nve interface nve auto-vlan-map show interfaces counters vlan	
Notes	• BGP controller mode must be set prior to using this command • For complete configuration, this VLAN needs to be created and a VXLAN source loopback needs to be added	

13.6.13 interface nve auto-vlan-map

	interface nve <nve> nve vni auto-vlan-map [base <base-number>] interface nve <nve> no nve vni auto-vlan-map Performs automatic mapping of all existing VLANs that are not manually mapped to VNI to a calculated VNI (Calculated VNI=base-number + VLAN). The no form of the command disables automatic VLAN mapping.	
Syntax Description	base-number	Range: 1-16773120 Default: 100000
Default	Disabled	
Configuration Mode	interface nve <nve>	
History	3.8.2200	

Example	`(config interface nve 1) # nve vni auto-vlan-map`
	`(config) # vlan 2-5`
	`(config) # show interfaces nve 1 detail`
	`----------------------------------------------------------------`
	`Vlan  VNI  Neigh Suppression  Mapping type`
	`----------------------------------------------------------------`
	`1  100001  Disabled  Auto`
	`2  100002  Disabled  Auto`
	`3  100003  Disabled  Auto`
	`4  100004  Disabled  Auto`
	`5  100005  Disabled  Auto`	

Related Commands	nve vni vlan interface nve disable nve vni	
Notes	• Base-number cannot be changed, user must unset auto-vlan-map and reconfigure it with a different base number • While auto-vlan-map is enabled, user cannot add manual mappings (only deletion of a manual mapping is allowed) • IPL VLAN will not be mapped to VNI.	

13.6.14 interface nve disable nve vni

	interface nve <nve> disable nve vni any vlan <vlan/vlan-range> interface nve <nve> no nve vni any vlan <vlan/vlan-range> Excludes a VLAN from the auto-vlan-map operation. The no form of the command deletes the exclusion.
Syntax Description	N/A
Default	Disabled
Configuration Mode	interface nve <nve>
History	3.8.2200
Example	`(config interface nve 1) # disable nve vni any vlan 5` `(config interface nve 1) # no nve vni any vlan 5`
Related Commands	interface nve auto-vlan-map
Notes	User can set/unset exclude VLANs while auto-vlan-map is enabled or disabled.

13.6.15 vxlan mlag-tunnel-ip

	vxlan mlag-tunnel-ip <mlag_ipv4_address> no vxlan mlag-tunnel-ip <mlag_ipv4_address> Configures the MLAG tunnel IP. The no form of the command unbinds VXLAN tunnel from the loopback interface.	
Syntax Description	mlag_ipv4_address	Valid MLAG IPv4 address
Default	N/A	
Configuration Mode	config interface nve	
History	3.8.1000	
Example	`switch (config interface nve 1) # vxlan mlag-tunnel-ip 1.2.3.4`	
Related Commands	protocol nve nve controller bgp	
Notes	BGP controller mode must be set prior to running this command	

13.6.16 vxlan source interface loopback

	vxlan source interface loopback <loopback-id> no vxlan source interface loopback <loopback-id> Binds VXLAN tunnel to a loopback interface. The no form of the command unbinds VXLAN tunnel from the loopback interface.	
Syntax Description	loopback-id	Loopback interface ID Range: 0-31
Default	N/A	
Configuration Mode	config interface nve	
History	3.6.3004	

Example	`switch (config interface nve 1) # vxlan source interface loopback 14`
Related Commands	protocol nve interface nve
Notes	• The configured loopback interface becomes the VXLAN tunnel endpoint (VTEP) • The configured loopback interface must be in the 'default ' VRF

13.6.17 shutdown

	shutdown no shutdown Disables VXLAN tunnel. The no form of the command enables VXLAN tunnel.
Syntax Description	N/A
Default	N/A
Configuration Mode	config interface nve
History	3.6.6102
Example	`switch (config interface nve 1) # shutdown`
Related Commands	protocol nve
Notes	

13.6.18 clear mac-address-table nve

	clear mac-address-table nve [remote] Clears locally-learned NVE MAC addresses.	
Syntax Description	remote	Clears remotely-learned NVE MAC addresses
Default	N/A	
Configuration Mode	config interface nve	
History	3.6.8008	
Example	`switch (config interface nve 1) # clear mac-address-table nve`	
Related Commands	protocol nve interface nve	
Notes		

13.6.19 clear nve counters

	clear nve counters Clears NVE counters.
Syntax Description	N/A
Default	N/A
Configuration Mode	config interface nve
History	3.6.3004

Example	switch (config interface nve 1) # clear nve counters
Related Commands	protocol nve interface nve
Notes	The command “clear counters all” also clears NVE counters

13.6.20 show interfaces nve

	show interfaces nve [<nve-id>] Displays information about NVE interfaces.	
Syntax Description	nve-id	NVE ID Range: 1-64
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
	3.8.1000	Updated example
	3.8.2200	Updated example and added auto-vlan-map status.
	3.9.0300	Updated example
Example	switch (config) # show interface nve 1 Admin state : enabled Source interface : loopback 1 Source interface ip : 192.168.1.1 Controller mode : BGP Mlag tunnel ip : not configured Effective tunnel ip : 192.168.1.1 Global neigh-suppression : Disabled Auto-vlan-map : Disabled Counters 1840 encapsulated (Tx) NVE packets 1970 decapsulated (Rx) NVE packets 0 dropped NVE-encapsulated packets 0 NVE-encapsulated packets with errors	
Related Commands		
Notes		

13.6.21 show interfaces nve detail

	show interfaces nve [<nve-id>] detail Displays all the VNI-VLAN mappings for this NVE interface.	
Syntax Description	nve-id	NVE ID Range: 1-64
Default	N/A	
Configuration Mode	Any command mode	
History	3.8.1000	
	3.8.2200	Added “Mapping type” to show whether VLAN to VNI mapping was done manually or by auto-vlan-map

	3.9.0300	Updated example
Example	switch (config)# show interfaces nve 1 detail Admin state : enabled Source interface : loopback 1 Source interface ip : 192.168.1.1 Controller mode : BGP Mlag tunnel ip : not configured Effective tunnel ip : 192.168.1.1 Global neigh-suppression : Disabled Auto-vlan-map : Disabled ---------- Vlan VNI Neigh Suppression Mapping Type ---------- 6 60 Disabled Manual 7 70 Disabled Manual 8 80 Disabled Manual 9 90 Disabled Manual	
Related Commands		
Notes		

13.6.22 show interfaces nve counters

	show interfaces nve <nve-id> counters Displays NVE counters.	
Syntax Description	nve-id	NVE ID Range: 1-64
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
	3.9.0500	Updated example
Example	switch (config) # show interface nve 1 counters encapsulated (Tx) NVE packets :0 decapsulated (Rx) NVE packets :0 dropped NVE-encapsulated packets :0 NVE-encapsulated packets with errors :0	
Related Commands		
Notes		

13.6.23 show interfaces counters vlan

	show interfaces nve <nve_id>counters vlan <vlan_value> Displays NVE counters per VLAN.	
Syntax Description	nve-id	NVE ID Range: 1-64
	vlan_value	VLAN value
Default	N/A	
Configuration Mode	Any command mode	
History	3.9.1000	

Example	switch (config) # show interfaces nve 1 counters vlan 5 Encapsulated (Tx) NVE packets: 1 Decapsulated (Rx) NVE packets: 1 Encapsulated (Tx) NVE bytes : 102 Decapsulated (Rx) NVE bytes : 152 switch (config) #
Related Commands	nve vni vlan
Notes	

13.6.24 show interfaces nve flood

	.show interfaces nve <nve-id> flood [vni <vni-id>] Displays remote VTEP endpoints configured for BUM (broadcast, unknown unicast, multicast) flooding.	
Syntax Description	nve-id	NVE ID Range: 1-64
	vni	Displays NVE flooding on specific VNI
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	

Example

```
switch (config) # show interface nve 1 flood

NVE Interface   Logical Switch         VNI ID     Flood IP Address
-------------   --------------         ------     ----------------
1               ls7777                 7777       1.2.3.5
```

Example (BGP controller mode)

```
switch (config) # show interfaces nve 1 flood

-----------------------------------------------------
NVE Interface VLAN ID VNI ID Flood IP Addresses
-----------------------------------------------------
1             6       60     192.168.1.2
1             7       70     193.168.1.1
                             193.168.1.2
```

Related Commands	
Notes	

13.6.25 show interfaces nve mac-address-table

	show interfaces nve <nve-id> mac-address-table [vni <vni-id>] Displays MAC address table of NVE interface.	
Syntax Description	nve-id	NVE ID Range: 1-64
	vni	Displays NVE flooding on specific VNI
Default	N/A	
Configuration Mode	Any command mode	

History	3.6.3004
Example	

```
switch (config) # show interface nve 1 mac-address-table

NVE Interface   Logical Switch         VNI ID     Mac Address         Address Type        Remote Endpoint
IP Address
-------------   --------------        ------     -----------         ------------
--------------------------
1               ls7777                7777       e4:1d:2d:a5:f2:0a   local learned       N/A
1               ls7777                7777       00:11:22:33:44:55   remote configured   1.2.3.5
```

Related Commands	
Notes	

13.6.26 show interfaces nve mac-address-table local learned unicast

	show interfaces nve <nve-id> mac-address-table local learned unicast [vni <vni-id>] Displays only the locally-learned unicast MAC addresses.	
Syntax Description	nve-id	NVE ID Range: 1-64
	vni	Displays NVE flooding on specific VNI
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Example		

```
switch (config) # show interface nve 1 mac-address-table local learned unicast

NVE Interface   Logical Switch         VNI ID     Mac Address         Address Type        Remote Endpoint
IP Address
-------------   --------------        ------     -----------         ------------
--------------------------
1               ls7777                7777       e7:3a:7e:a5:f2:1a   local learned       N/A
```

Related Commands	
Notes	

13.6.27 show interfaces nve mac-address-table remote configured multicast

	show interfaces nve <nve-id> mac-address-table remote configured multicast [vni <vni-id>] Displays only remotely-configured BUM addresses.	
Syntax Description	nve-id	NVE ID Range: 1-64
	vni	Displays NVE flooding on specific VNI
Default	N/A	
Configuration Mode	Any command mode	

History	3.6.3004
Example	

```
switch (config) # show interface nve 1 mac-address-table remote configured multicast

NVE Interface   Logical Switch         VNI ID     Mac Address         Address Type        Remote Endpoint
IP Address
-------------   --------------        ------     -----------         ------------
--------------------------
1               ls7777                7777       00:11:22:33:44:55   remote configured   1.2.3.5
```

Related Commands	
Notes	

13.6.28 show interfaces nve peers

	show interfaces nve <nve-id> peers [vni <vni-id>] Displays all remote VTEPs.	
Syntax Description	nve-id	NVE ID Range: 1-64
	vni	Displays NVE flooding on specific VNI
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
	3.8.2200	Added output of the command while running NVE BGP controller mode
Example		

```
switch (config) # show interfaces nve 1 peers

-----------------------------------------------------
NVE Interface  Logical Switch  VNI ID   Peer IP Address
-----------------------------------------------------
1              bridge          10080    1.1.1.1
1              bridge          10080    1.1.1.2
```

When running in NVE BGP controller mode:

```
switch (config) # show interfaces nve 1  peers
-------------------------------------------------
NVE Interface  VLAN ID  VNI ID   Peer IP Address
-------------------------------------------------
1              5        50       192.168.1.1
1              6        60       192.168.1.1
```

Related Commands	
Notes	

13.6.29 ovs ovsdb server

	ovs ovsdb server no ovs ovsdb server Runs OVSDB-server process and unhides OVS commands. The no form of the command deactivates OVSDB-server process and hides OVS commands.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.3004
Example	`switch (config) # ovs ovsdb server`
Related Commands	
Notes	OVSDB server runs when “protocol openflow” or “protocol nve” are enabled, even when not enabled using this command

13.6.30 ovs ovsdb manager remote

	ovs ovsdb manager remote {tcp | ssl} ip-address <ip-address> port <tcp-port> no ovs ovsdb manager remote {tcp | ssl} ip-address <ip-address> port <tcp-port> Configures OVSDB to actively connect to a remote manager at a given IP address and TCP port, using either TCP or SSL. The no form of the command disconnects OVSDB from a remote manager.	
Syntax Description	SSL	Connect with TCP protocol
	TCP	Connect with SSL protocol
	ip-address	IP address of remote manager
Default	N/A	
Configuration Mode	config	
History	3.6.3004	
Example	`switch (config) # ovs ovsdb manager remote tcp ip-address 10.10.10.10 port 20`	
Related Commands	ovs ovsdb server	
Notes		

13.6.31 ovs ovsdb server listen

	ovs ovsdb server listen {tcp | ssl} port <tcp-port> local ip-address <ip-address> no ovs ovsdb server listen {tcp | ssl} port <tcp-port> local ip-address <ip-address> Configures OVSDB to listen at a given port of an interface with a given (local) IP address. The no form of the command disconnects OVSDB from a remote manager.	
Syntax Description	SSL	Connect with TCP protocol
	TCP	Connect with SSL protocol
	ip-address	IP address of a given port
Default	N/A	
Configuration Mode	config	
History	3.6.3004	
Example	`switch (config) # ovs ovsdb server listen tcp port 20 local ip-address 20.20.20.20`	
Related Commands	ovs ovsdb server	

Notes	

13.6.32 ovs logging level

	ovs {ovsdb | vswitchd | vtep} logging level {dbg | emer | err | info | off | warn} Configures OVS logging levels for OVS related processes.	
Syntax Description	ovsdb | vswitchd | vtep	OVS-related processes
	dbg | emer | err | info | off | warn	Logging level severity
Default	N/A	
Configuration Mode	config	
History	3.8.1100	
Example	`switch (config) # ovs ovsdb logging level err` `switch (config) # ovs ovsdb vswitchd level warn`	
Related Commands		
Notes		

13.6.33 show ovs

	show ovs Displays OVS information.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.8.1100

Example	`switch (config) # show ovs`
	
	`Logging level:`
	`  ovsdb   : info`
	`  vswitchd: info`
	`  vtep    : warn`

Related Commands	
Notes	

14 Ethernet VPN (EVPN)

14.1 Overview

Many data centers today are moving from legacy Layer 2 (L2) designs to modern Layer 3 (L3) web-scale IT architectures. L3 designs simplify troubleshooting, provide clear upgrade strategies, support multi-vendor environments, and dramatically reduce the size of failure domains.

General Data Center Network with EVPN

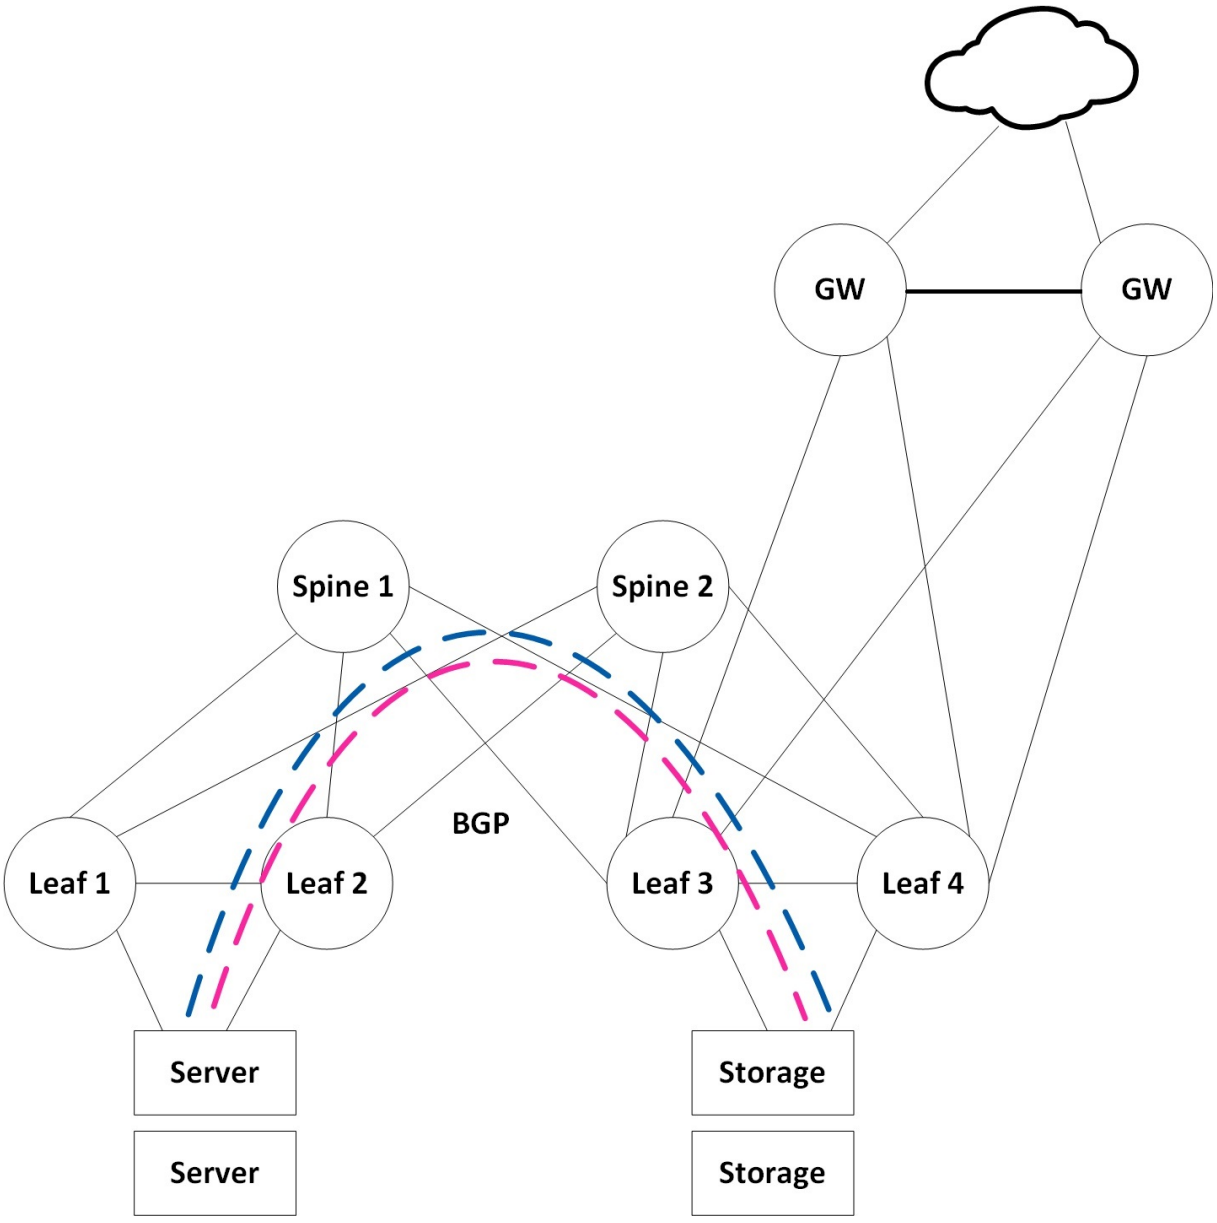

However, many applications and storage appliances still require layer 2 adjacency. VXLAN tunnels can satisfy this L2 adjacency requirement, and EVPN serves as a standard for scale-out L2 Ethernet

fabrics. VXLAN can virtualize the data center network, enabling layer 2 segments to be extended over an IP core (the underlay). EVPN is the control plane for modern VXLAN deployments, allowing VTEPs to discover each other via EVPN and exchange reachability information such as MAC and IPs across racks.

ARP suppression is used to reduce the amount of broadcast packets crossing the extended L2 domain. BGP is the underlay routing protocol serving as the transport layer for the overlay VXLAN.

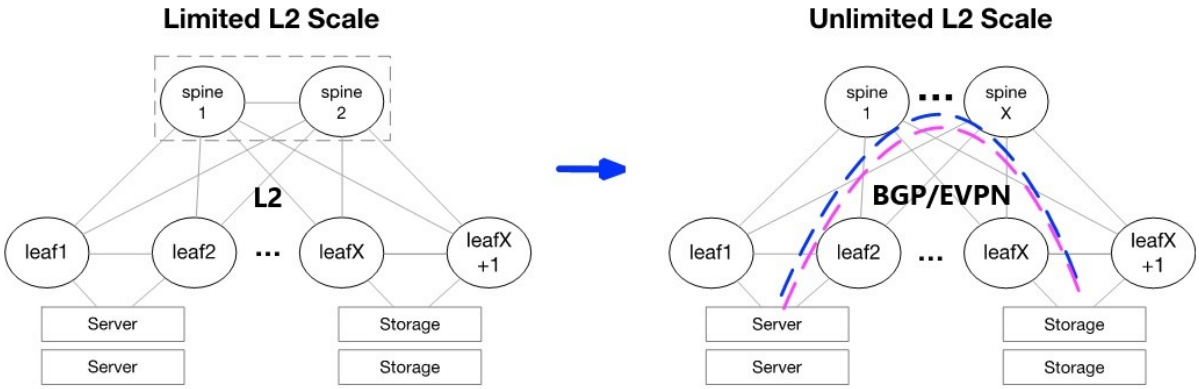

14.2 Example of How To Configure EVPN

The configuration flow will be described using the setup illustrated below and over leaf3.

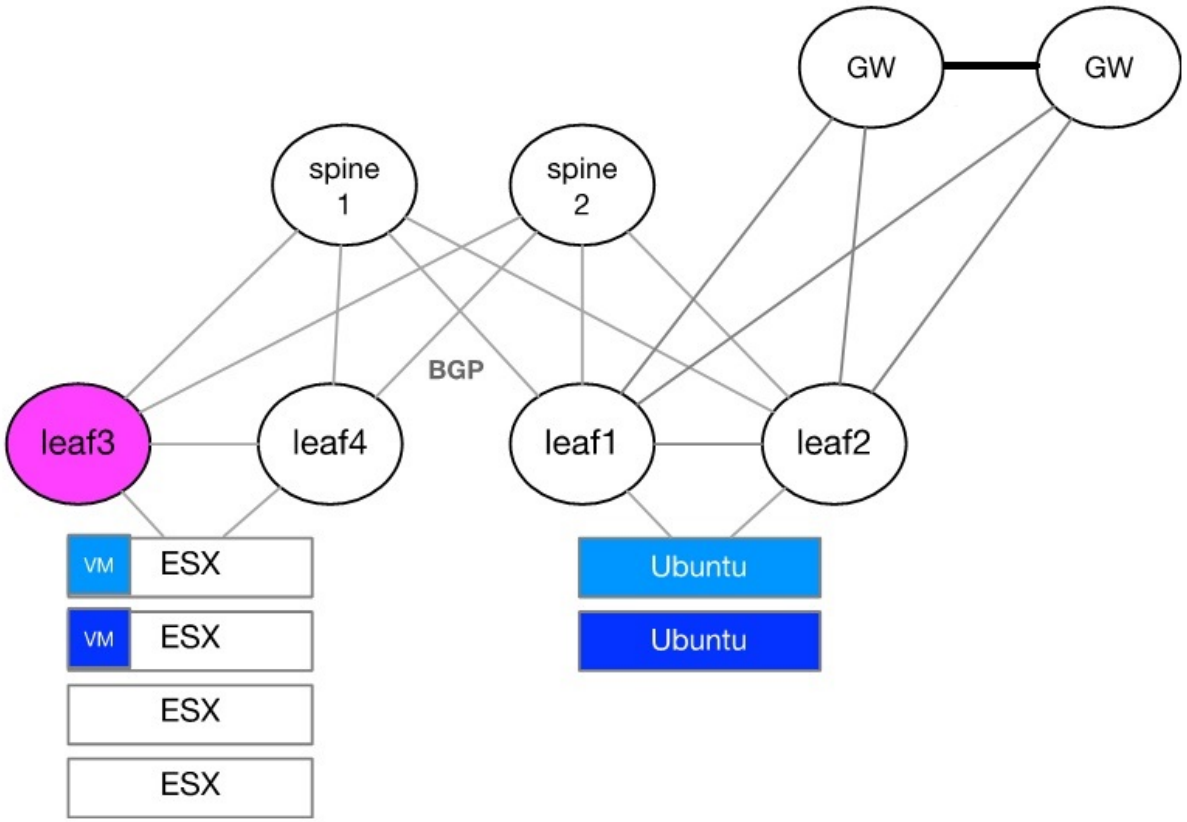

14.2.1 Layer 2 Configuration, MLAG, and VLANs

MLAG between leaf3 and leaf4

```
lacp
dcb priority-flow-control enable force
protocol mlag
interface port-channel 1
interface ethernet 1/1 channel-group 1 mode active
interface port-channel 1 dcb priority-flow-control mode on force
interface mlag-port-channel 7-8 no shutdown
interface ethernet 1/31 mlag-channel-group 7 mode active
interface ethernet 1/32 mlag-channel-group 8 mode active
vlan 4094
ip routing vrf default
interface vlan 4094
interface vlan 4094 ip address 10.10.10.1/30 primary
interface vlan 4094 mtu 9216
mlag-vip mlag-pair-1 ip 192.168.1.1 /24 force
interface port-channel 1 ipl 1
interface vlan 4094 ipl 1 peer-address 10.10.10.2
no mlag shutdown
```

Layer 2 Ports

- In our setup we use VLAN 6 as the native VLAN, and VLAN 10 as the Tagged VLAN.
- We use LACP Bond on our servers, and using them we set LACP on the Switch MPOs.
- PXE boot is required to set our MPOs to "lacp-individual enable"

```
interface mlag-port-channel 7-8
interface ethernet 1/7 mlag-channel-group 7 mode active
interface ethernet 1/8 mlag-channel-group 8 mode active
interface mlag-port-channel 7-8 mtu 9216 force
interface mlag-port-channel 7 switchport mode hybrid
interface mlag-port-channel 8 switchport mode hybrid
interface mlag-port-channel 7-8 no shutdown
lacp
interface mlag-port-channel 7-8 lacp-individual enable force
vlan 6
vlan 10
interface mlag-port-channel 7 switchport access vlan 6
interface mlag-port-channel 8 switchport access vlan 6
interface mlag-port-channel 7 switchport hybrid allowed-vlan 10
interface mlag-port-channel 8 switchport hybrid allowed-vlan 10
```

14.2.2 Layer 3 Configuration

Layer 3 Interfaces

- Since we use VXLAN, we will set all of our L3 interfaces to support a maximum MTU of 9216. The servers' MTU should be set to below the maximum fabric MTU to allow space for the additional headers of the VXLAN. The VXLAN encapsulation header adds 50 bytes to the overall size of an Ethernet frame.
- Router ports serve as uplinks.
- Loopback for VTEP source is unique per leaf switch.

```
interface ethernet 1/28 no switchport force
interface ethernet 1/29 no switchport force
interface ethernet 1/28 mtu 9216 force
interface ethernet 1/29 mtu 9216 force
interface loopback 1
interface ethernet 1/28 ip address 100.100.100.1/30 primary
interface ethernet 1/29 ip address 100.100.100.5/30 primary
interface loopback 1 ip address 1.1.1.1/32 primary
```

VXLAN Tunnels Configuration

NVE represents a VTEP. We will use a single VTEP with multiple VNIs.

```
protocol nve
interface nve 1
interface nve 1 vxlan source interface loopback 1
interface nve 1 nve controller bgp
interface nve 1 vxlan mlag-tunnel-ip 100.0.0.1
interface nve 1 nve vni 10010 vlan 10
interface nve 1 nve vni 10060 vlan 6
```

Note that "vxlan mlag-tunnel-ip" is used to configure MLAG with VXLAN. This way other VTEPs will see the MLAG pair as a single entity (for this reason, the "mlag-tunnel-ip" setting should be unique per MLAG pair). As long as the MLAG is up, both switches will use the same IP as the VTEP source. If MLAG state changes to Split Brain (IPL is down but mgmt0 interface is up), the standby switch will use its local loopback for the advertisements; this will prevent impacting traffic from stand-alone ports by the Split Brain scenario.

The only command needed to add more VNIs to a switch is:

```
interface nve 1 nve vni 10020 vlan 20
```

ARP Suppression

Traditional L2 network broadcast traffic generated by ARP requests overloads the network. Using ARP suppression with VXLAN enables suppressing these messages at the leaf layer. Let's consider the example setup that is illustrated below.

The support for gratuitous ARP in EVPN has been added, also when ARP Suppression is enabled. The feature allows generating GARP packets on the egress VTEP only when neighbor-suppression is enabled on both VTEPs in the chain (ingress and egress). The suppression should be enabled on interface NVE or on a particular VLAN of the VTEP.

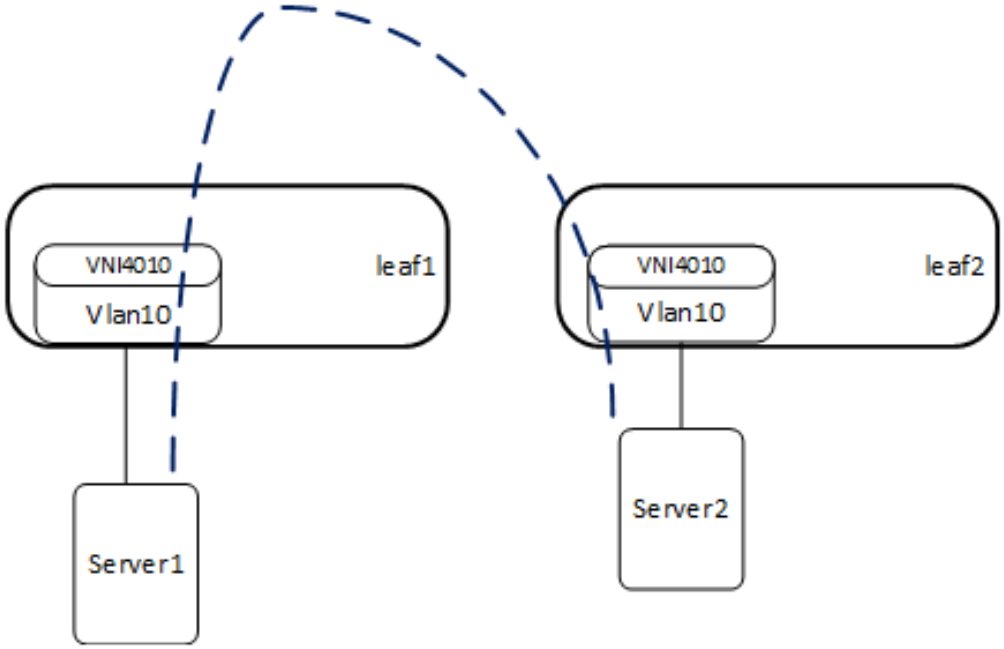

- The first time Server2 communicates, it sends an ARP request.
- Leaf2 learns its MAC and IP, and sends an EVPN update containing the IP and MAC on the corresponding VNI4010.
- Leaf1 learns the IP and MAC of Server2 on VNI4010.
- When Server1 sends an ARP request to Server2, leaf1 replies to the ARP request as it has all of the details.

- The result is that broadcasts to all leafs that are part of VNI4010 are suppressed.

```
interface nve 1 nve neigh-suppression
interface vlan 6
interface vlan 10
```

EVPN Neighbor-Suppression

	EVPN neighbor-suppression	
	enabled	disabled
IPv4 Normal ARP	suppressed	flooded
IPv4 GARP	suppressed*	flooded
IPv6 Neighbor Discovery (equivalent to IPv4 ARP)	suppressed	flooded
IPv6 Unsolicited Neighbor Advertisement (equivalent to IPv4 GARP)	flooded	flooded

* the GARP (Gratuitous ARP) packet will reach the destination endpoint despite neighbor suppression

Since IPv4 GARP is processed locally on ingress VTEP and then only BGP update propagated through EVPN network there are several limitations related to scale and performance. The following limitations vary based on the CPU type and current switch load. Switches with higher performance will have better results. Below are the minimum performance expected.

- Ingress VTEP: max 1000 frames per second of ingress GARP
- Egress VTEP: at least 100 fps for GARP generation

14.2.3 BGP and EVPN Configuration

> The examples below use eBGP. Nevertheless, iBGP can be used as well.

Now we will configure our L3 underlay using eBGP as the underlay protocol. The Autonomous System (AS) design that we use as an example represents common designs of eBGP running over leaf/spine data centers. Specifically, each of the leaf switches will be in a separate AS, and the spine layer will be in the same AS layer.

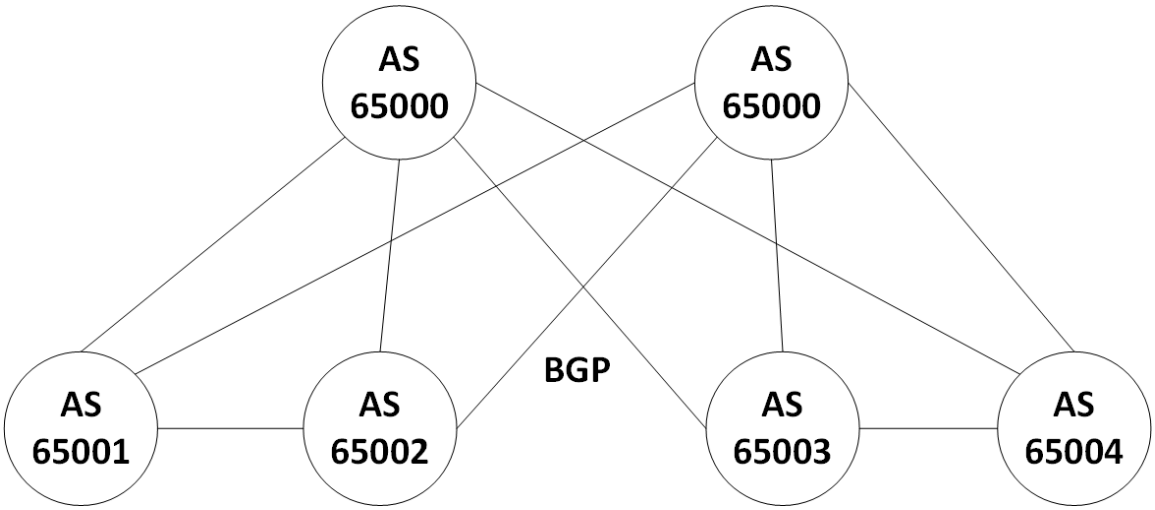

BGP

```
protocol bgp
router bgp 65001 vrf default
router bgp 65001 vrf default bgp fast-external-fallover
router bgp 65001 vrf default maximum-paths 32
router bgp 65001 vrf default bestpath as-path multipath-relax force
router bgp 65001 vrf default neighbor 10.10.10.2 remote-as 65002
router bgp 65001 vrf default neighbor 100.100.100.1 remote-as 65000
router bgp 65001 vrf default neighbor 100.100.100.5 remote-as 65000
router bgp 65001 vrf default network 1.1.1.1 /32
router bgp 65001 vrf default network 100.0.0.1 /32
```

Note: It is necessary to advertise both the local loopback network and the mlag-tunnel-ip network.

EVPN Address Family

In the following code, we create a peer group that contains all of the EVPN configuration and attach it to our L3 interfaces.

```
router bgp 65001 vrf default neighbor evpn peer-group
router bgp 65001 vrf default neighbor evpn send-community
router bgp 65001 vrf default neighbor evpn send-community extended
router bgp 65001 vrf default address-family l2vpn-evpn neighbor evpn next-hop-unchanged
router bgp 65001 vrf default address-family l2vpn-evpn neighbor evpn activate
router bgp 65001 vrf default address-family l2vpn-evpn vni auto-create
router bgp 65001 vrf default neighbor 10.10.10.1 peer-group evpn
router bgp 65001 vrf default neighbor 100.100.100.1 peer-group evpn
router bgp 65001 vrf default neighbor 100.100.100.5 peer-group evpn
```

14.2.4 Spine Configuration

Each spine has a unique loopback address that we use to represent its Router-ID.

```
ip routing vrf default
interface ethernet 1/1-1/4 no switchport force
interface ethernet 1/1-1/4 mtu 9216 force
interface loopback 1
interface ethernet 1/1 ip address 100.100.100.2/30 primary
interface ethernet 1/2 ip address 100.100.100.6/30 primary
interface ethernet 1/3 ip address 100.100.100.10/30 primary
interface ethernet 1/4 ip address 100.100.100.14/30 primary
interface loopback 1 ip address 1.1.1.5/32 primary
```

```
protocol bgp
router bgp 65000 vrf default
router bgp 65000 vrf default bgp fast-external-fallover
router bgp 65000 vrf default maximum-paths 32
router bgp 65000 vrf default bestpath as-path multipath-relax force
router bgp 65000 vrf default neighbor 100.100.100.1 remote-as 65001
```

```
router bgp 65000 vrf default neighbor 100.100.100.5 remote-as 65002
router bgp 65000 vrf default neighbor 100.100.100.9 remote-as 65003
router bgp 65000 vrf default neighbor 100.100.100.13 remote-as 65004
router bgp 65000 vrf default neighbor evpn peer-group
router bgp 65000 vrf default neighbor evpn send-community
router bgp 65000 vrf default neighbor evpn send-community extended
router bgp 65000 vrf default address-family l2vpn-evpn neighbor evpn next-hop-unchanged
router bgp 65000 vrf default address-family l2vpn-evpn neighbor evpn activate
router bgp 65000 vrf default neighbor 100.100.100.1 peer-group evpn
router bgp 65000 vrf default neighbor 100.100.100.5 peer-group evpn
router bgp 65000 vrf default neighbor 100.100.100.9 peer-group evpn
router bgp 65000 vrf default neighbor 100.100.100.13 peer-group evpn
router bgp 65000 vrf default network 1.1.1.5 /32
```

14.3 Traffic Behavior During Failures

Server Link Failure

Traffic forwarding during a failure follows standard MLAG behavior. If a link of the server fails, traffic will be forwarded across one of the remaining active links.

With reference to the illustration below: If traffic is received on leaf3 due to the ECMP hash of the spine, leaf3 will decapsulate the frame. And based on its local MAC table, leaf3 will also switch the frame across the peer link for forwarding to Server via leaf4.

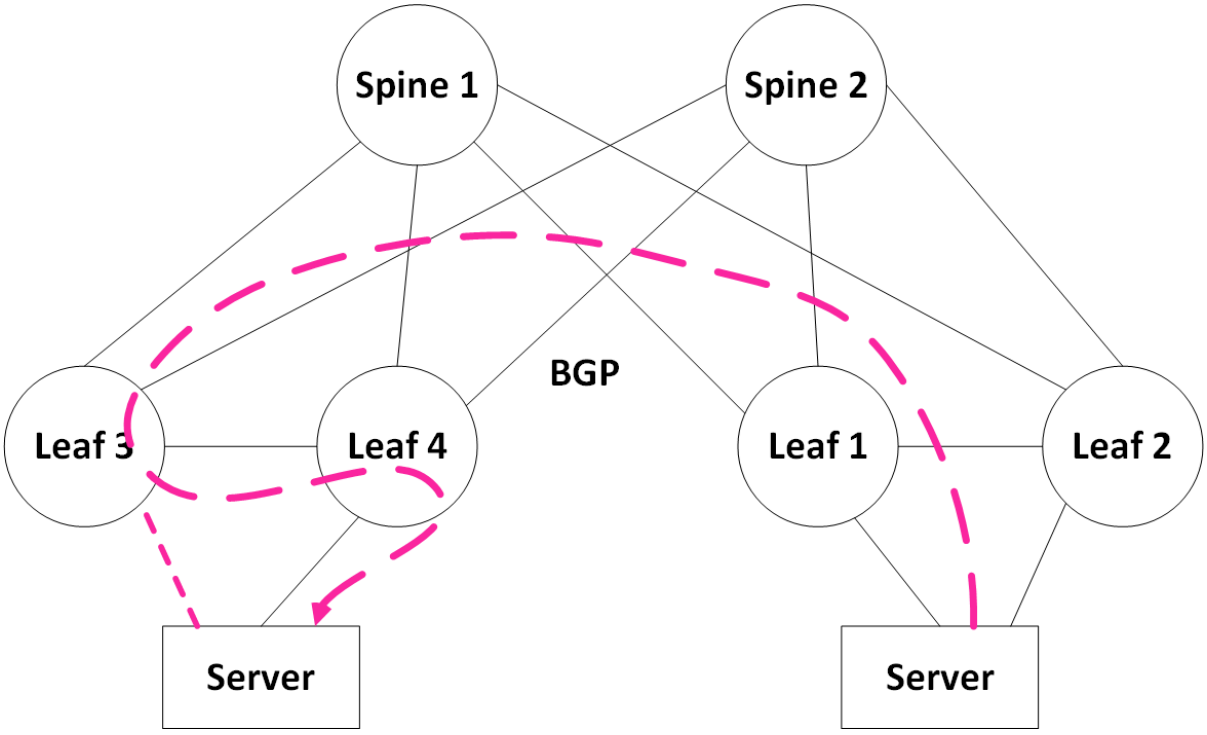

Uplink Failure

To cover rare cases such as losing all of the uplinks on one of the MLAG peers, we enable BGP over the IPL. This way, traffic coming from the servers towards that leaf can still be routed towards the

remote servers.

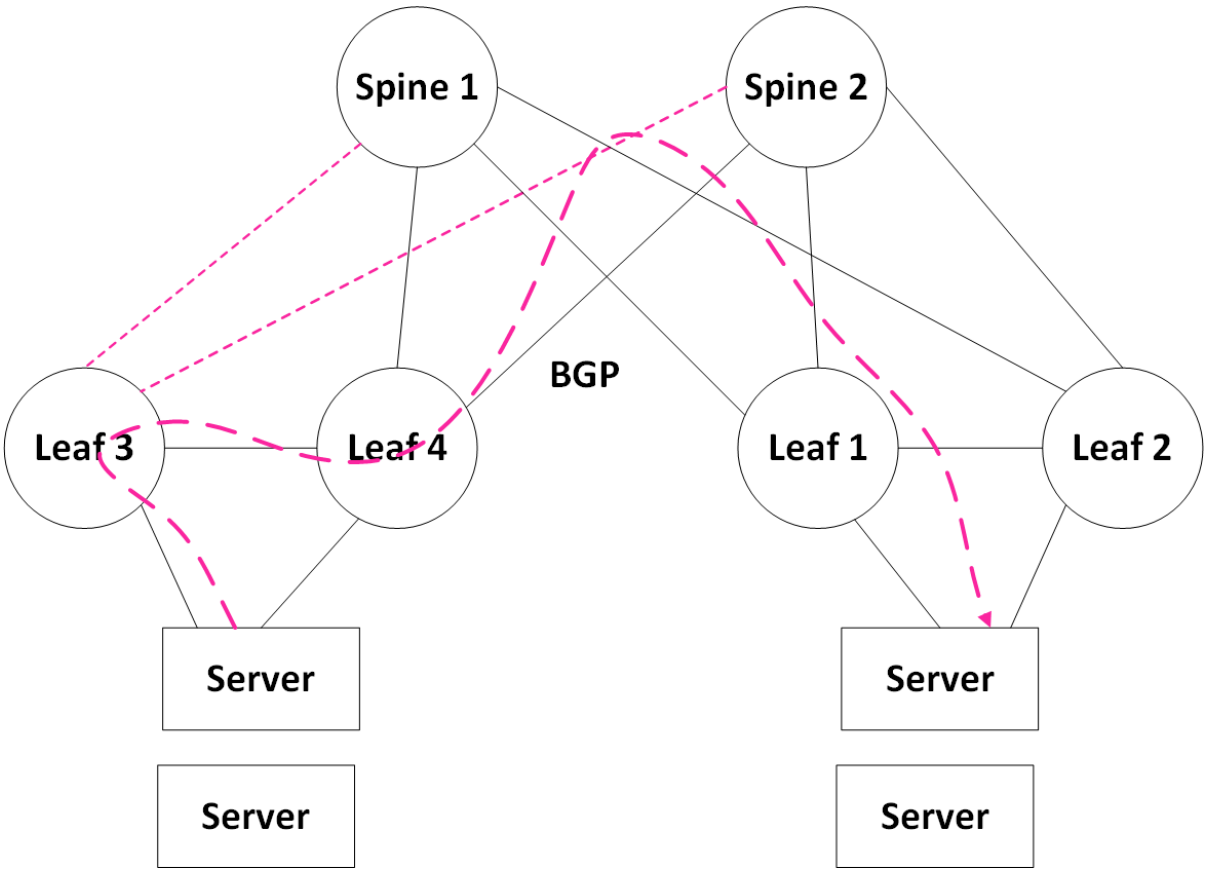

Note: Traffic coming towards the servers connected to leaf4 from the spine will always be terminated on leaf4 and sent directly to the servers without passing over the IPL.

14.4 EVPN Troubleshooting

14.4.1 show interface nve 1

Display the configured VTEP on a network device participating in BGP EVPN.

```
Interface NVE 1 status:
  Admin state             : Enabled
  Source interface        : loopback 1
  Source interface ip     : 4.4.4.4
  Controller mode         : BGP
  MLAG tunnel ip          : 8.8.8.8
  Effective tunnel ip     : 8.8.8.8
  Global neigh-suppression: Disabled
  Auto-vlan-map           : Enabled
  Auto-vlan-map base      : 100000

  Counters:
    encapsulated (Tx) NVE packets       : 0
    decapsulated (Rx) NVE packets       : 0
    dropped NVE-encapsulated packets    : 0
    NVE-encapsulated packets with errors: 0
```

14.4.2 show interface nve 1 detail

Display the configured VNIs on a network device participating in BGP EVPN.

```
Admin state             : Enabled
Source interface        : loopback 1
Source interface ip     : 4.4.4.4
Controller mode         : BGP
MLAG tunnel ip          : 8.8.8.8
Effective tunnel ip     : 8.8.8.8
Global neigh-suppression: Disabled
Auto-vlan-map           : Enabled
Auto-vlan-map base      : 100000

-------------------------------------------------------------
Vlan         VNI          Neigh Suppression     Mapping type
-------------------------------------------------------------
1            100001       Disabled              Auto
10           10           Disabled              Manual
20           20           Disabled              Manual
30           30           Disabled              Manual
4000         Excluded     N/A                   N/A
```

14.4.3 show ip bgp evpn summary

Display the BGP peers participating in the layer 2 EVPN address-family and their states.

```
VRF name                  : default
BGP router identifier     : 1.1.1.1
local AS number           : 101
BGP table version         : 2176
Main routing table version: 2176
IPV4 Prefixes             : 12
IPV6 Prefixes             : 0
L2VPN EVPN Prefixes       : 9

-------------------------------------------------------------------------------------------------------------------
Neighbor          V    AS           MsgRcvd   MsgSent   TblVer    InQ    OutQ   Up/Down       State/PfxRcd
-------------------------------------------------------------------------------------------------------------------
10.10.10.2        4    102          2320      2539      2176      0      0      0:00:46:52    ESTABLISHED/5
192.168.14.4      4    104          2112      3159      2176      0      0      0:00:57:56    ESTABLISHED/4
```

14.4.4 show ip bgp evpn

Display all EVPN routes, both local and remote. The routes displayed here are based on RD as they are across VNIs.

```
show ip bgp evpn

-------------------------------------------------------------------------------------------------------------------
-----------------
RD                    Type              Data                    Next Hop          Metric          LocPrf       Weight
Path
-------------------------------------------------------------------------------------------------------------------
-----------------
1.1.1.1:10            mac-ip            00:00:01:11:22:33       1.1.1.1           0               100          0
104 101 ?
9.9.9.9:30            mac-ip            00:10:94:00:00:02       9.9.9.9           0               100          0
104 102 ?
9.9.9.9:30            mac-ip            00:10:94:00:00:03       9.9.9.9           0               100          0
104 102 ?
1.1.1.1:10            imet              1.1.1.1                 1.1.1.1           0               100          0
104 101 ?
3.3.3.3:10            imet              3.3.3.3                 3.3.3.3           0               100          0
          ?
3.3.3.3:30            imet              3.3.3.3                 3.3.3.3           0               100          0
          ?
3.3.3.3:123           imet              3.3.3.3                 3.3.3.3           0               100          0
          ?
9.9.9.9:10            imet              9.9.9.9                 9.9.9.9           0               100          0
104 102 ?
9.9.9.9:30            imet              9.9.9.9                 9.9.9.9           0               100          0
104 102 ?
```

14.4.5 show ip bgp evpn vni 10060

Display the EVPN information for a specific VNI in detail.

```
show ip bgp evpn vni 10060
```

```
------------------------------------------------------------------------------------------------------------------------
--------------------------
RD                      Type              Data                        Next Hop         Metric          LocPrf
Weight          Path
------------------------------------------------------------------------------------------------------------------------
--------------------------
1.1.1.1:321             mac-ip            00:00:01:11:22:33           1.1.1.1          0               100
0           104 101 ?
1.1.1.1:321             mac-ip            00:10:00:00:00:05           1.1.1.1          0               100
0           104 101 ?
1.1.1.1:321             mac-ip            00:10:33:01:7d:2a           1.1.1.1          0               100
0           104 101 ?
1.1.1.1:321             mac-ip            00:10:66:02:fa:54           1.1.1.1          0               100
0           104 101 ?
1.1.1.1:321             mac-ip            00:10:88:06:a7:33           1.1.1.1          0               100
0           104 101 ?
1.1.1.1:321             mac-ip            00:10:cc:05:f4:a8           1.1.1.1          0               100
0           104 101 ?
9.9.9.9:321             mac-ip            00:10:94:00:00:02           9.9.9.9          0               100
0           104 102 ?
9.9.9.9:321             mac-ip            00:10:94:00:00:03           9.9.9.9          0               100
0           104 102 ?
```

14.4.6 show ip bgp evpn with multiple filters

Display the EVPN information for a specific VNI in detail, selecting different filters

```
switch (config) # show ip bgp evpn vni 1000 route-type mac-ip

------------------------------------------------------------------------------------------------------------------------
--------------------
RD                      Type              Data                          Next Hop         Metric      LocPrf
Weight      Path
------------------------------------------------------------------------------------------------------------------------
--------------------
2.3.4.5:5               mac-ip            00:bb:cc:dd:ee:ff             2.3.4.5          0           100         0
            ?

switch (config) # show ip bgp evpn vni 1000 route-type mac-ip detail

1 paths for mac-ip 00:bb:cc:dd:ee:ff   Route Distinguisher: 2.3.4.5:5:
  route:
    next hop          : 2.3.4.5
    neighbor ip       : 1.1.1.2
    router id         : 2.3.4.5
    metric            : 0
    weight            : 0
    local pref        : 100
    origin            : incomplete
    Extended Community: 100:268436456(Route-Target-AS)
    Extended Community: tunnelTypeVxlan(TunnelEncap)
    flags             : valid, best
    esi               : 00:00:00:00:00:00:00:00:00:00
    vni               : 1000
    path              :
    ethernet tag id   :
```

14.4.7 show mac-address-table

Display all local and remote MAC addresses.

```
------------------------------------------------------------
Vlan    Mac Address          Type          Port\Next Hop
------------------------------------------------------------
10      00:00:01:11:22:33    Static        9.9.9.9(nve1)
10      00:00:01:55:A4:25    Static        1.1.1.1(nve1)
10      00:10:00:00:0A:67    Dynamic       Eth1/10
10      00:10:44:03:51:01    Dynamic       Eth1/10
10      00:10:88:06:A2:02    Dynamic       Eth1/10
10      00:10:AA:07:0F:B1    Dynamic       Eth1/10
30      00:10:00:00:05:29    Dynamic       1.1.1.1(nve1)
30      00:10:00:00:0A:52    Dynamic       1.1.1.1(nve1)
123     00:10:00:00:0A:5B    Dynamic       9.9.9.9(nve1)
123     00:10:44:03:51:0E    Dynamic       9.9.9.9(nve1)
123     00:10:88:06:A2:1C    Dynamic       9.9.9.9(nve1)


Number of unicast(local): 4
Number of NVE:        7
```

14.4.8 show ip arp

Display all local and remote neighbors (ARP entries), this command is only relevant when arp-suppression is enabled.

```
VRF Name default:
Total number of entries: 13


--------------------------------------------------------------------------------
Address              Type            Hardware Address          Interface
--------------------------------------------------------------------------------
10.209.20.1          Dynamic ETH     00:00:5E:00:01:01         mgmt0
10.209.20.57         Dynamic ETH     90:B1:1C:04:11:8D         mgmt0
10.209.20.58         Dynamic ETH     90:B1:1C:04:11:C1         mgmt0
10.209.20.67         Dynamic ETH     90:B1:1C:03:57:09         mgmt0
151.151.10.1         Dynamic ETH     98:03:9B:A2:BF:80         mgmt0
136.6.166.105        Dynamic ETH     00:00:66:02:FB:0C         vlan 10
136.6.162.102        Dynamic EVPN    00:00:00:00:05:58         vlan 123
136.6.162.114        Dynamic EVPN    00:00:01:00:00:02         vlan 123
172.3.12.4           Static EVPN     00:11:22:33:44:55         vlan 123
136.6.165.153        Dynamic EVPN    00:00:44:03:51:30         vlan 30
136.6.166.99         Dynamic EVPN    00:00:01:00:00:02         vlan 30
204.5.245.253        Dynamic EVPN    00:00:22:01:A8:98         vlan 30
192.168.34.4         Dynamic ETH     24:8A:07:F4:FF:48         eth 1/15
```

14.5 EVPN Data Center Interconnect (DCI)

14.5.1 Layer 2 DCI Connection

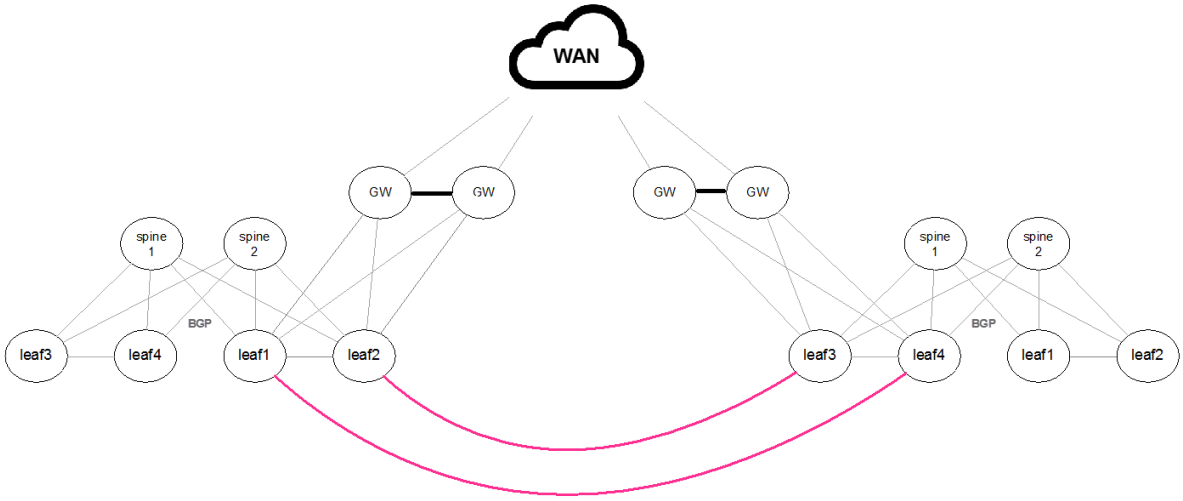

Regular BGP/EVPN Configuration is required since the connection between the sites is L2 based.

14.5.2 Layer 3 Routes WAN

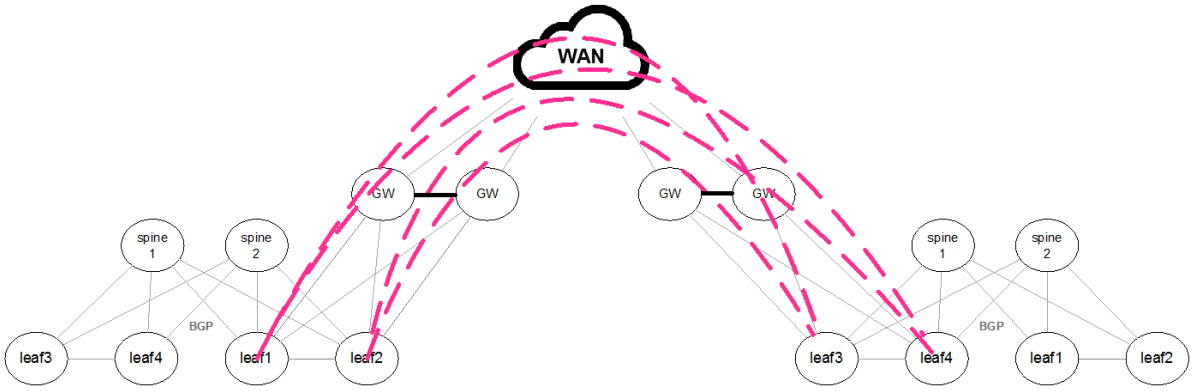

As the WAN transport layer does not support the EVPN/BGP address family, a remote BGP/EVPN connection should be set between each of the local leafs and the remote leafs. To allow this connection BGP should be set to multi-hop mode.

```
router bgp 65004 neighbor 100.100.100.5 ebgp-multihop 254
```

14.6 EVPN Centralized L3 Gateway

In centralized L3 gateway, a specific VTEP can be configured to act as the default gateway for all the hosts in a particular subnet throughout the EVPN network. It is possible to provision an MLAG pair in active-active mode as the default gateway. The VTEP will perform a routing to the destination host together with VxLAN ingress and egress bridging.

14.6.1 Configuration Example of EVPN Centralized Gateway

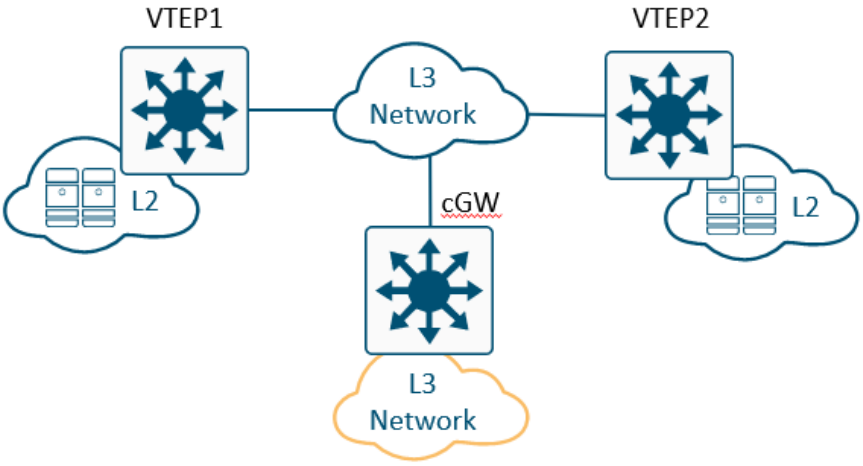

Run the following:

```
ip routing vrf default
protocol nve
interface nve 1
interface nve 1 vxlan source interface loopback 1
```

```
interface nve 1 nve controller bgp
interface nve 1 nve vni auto-vlan-map base 100000
```

- The underlay

```
protocol bgp
router bgp 1 vrf default
router bgp 1 vrf default router-id 200.0.1.1 force
router bgp 1 vrf default neighbor 1.1.1.2 remote-as 1
router bgp 1 vrf default address-family l2vpn-evpn neighbor 1.1.1.2 send-community
router bgp 1 vrf default address-family l2vpn-evpn neighbor 1.1.1.2 send-community extended
router bgp 1 vrf default address-family l2vpn-evpn neighbor 1.1.1.2 next-hop-unchanged
router bgp 1 vrf default address-family l2vpn-evpn neighbor 1.1.1.2 activate
router bgp 1 vrf default redistribute connected
router bgp 1 vrf default address-family l2vpn-evpn vni auto-create
```

- The overlay

```
vlan 10
interface vlan 10 ip address 192.168.1.1 /24
```

VTEP Key Outputs

```
VTEP1 switch (config) # show ip bgp evpn detail

1 paths for mac-ip b8:59:9f:a7:0f:88 192.168.1.1 Route Distinguisher: 200.0.1.1:10:
  route:
    next hop           : 1.1.1.1
    neighbor ip        : 2.2.2.2
    router id          : 200.0.2.1
    metric             : 0
    weight             : 0
    local pref         : 100
    origin             : incomplete
    Extended Community: 1:269205466(Route-Target-AS)
    Extended Community: tunnelTypeVxlan(TunnelEncap)
    Extended Community: defaultGateway
    flags              : valid, best
    esi                : 00:00:00:00:00:00:00:00:00:00
    vni                : 100010
    path               : 1
    ethernet tag       : 0


VTEP1 switch (config) # show ip arp

Flags:
  G: EVPN Default GW

VRF Name default:
  Total number of entries: 2

  ------------------------------------------------------------------------------------------------
  Address               Type             Flags      Hardware Address          Interface
  ------------------------------------------------------------------------------------------------
  192.168.1.1           Dynamic EVPN     G          B8:59:9F:A7:0F:88         vlan 10
```

14.6.2 Configuration Example of MLAG EVPN Centralized Gateway

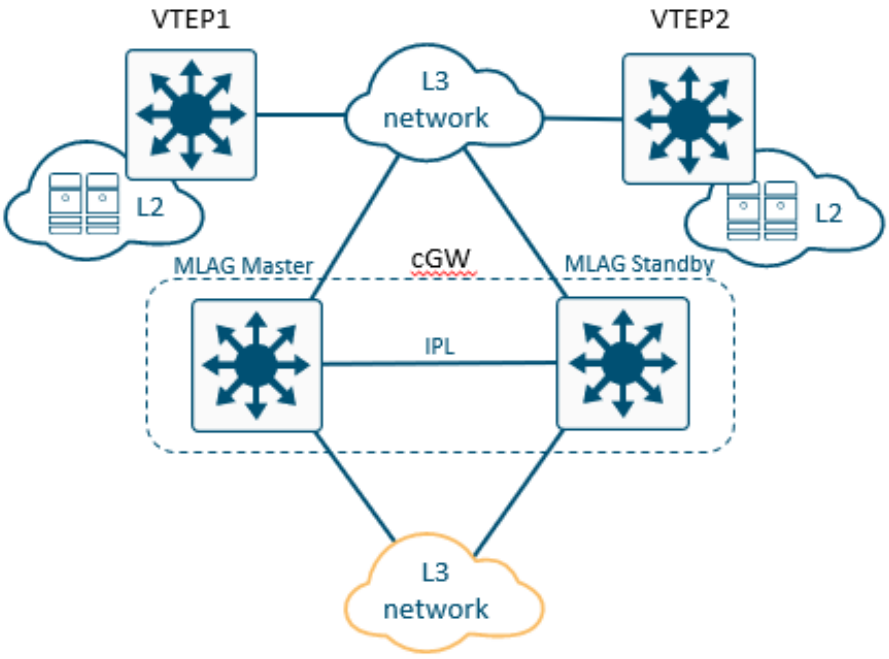

1. Configure the MLAG Master in the following way

```
ip routing vrf default
protocol nve
interface nve 1
interface nve 1 vxlan source interface loopback 1
interface nve 1 nve controller bgp
interface nve 1 vxlan mlag-tunnel-ip 5.5.5.5
interface nve 1 nve vni auto-vlan-map base 100000

interface mlag-port-channel 1 switchport access vlan 10
interface vlan 10 ip address 10.0.0.251/24 primary
interface vlan 4094 ip address 2.2.2.1/24 primary

protocol magp
interface vlan 10 magp 10
interface vlan 10 magp 10 ip virtual-router address 10.0.0.1
interface vlan 10 magp 10 ip virtual-router mac-address 00:00:5e:00:00:10

mlag-vip MLAG-1 ip 11.11.11.11 /24 force
interface port-channel 1 ipl 1
interface vlan 4094 ipl 1 peer-address 2.2.2.2

protocol bgp
router bgp 1 vrf default
router bgp 1 router-id 200.0.0.1 force
router bgp 1 neighbor 1.1.1.2 remote-as 1
router bgp 1 neighbor 2.2.2.2 remote-as 1
router bgp 1 address-family l2vpn-evpn neighbor 1.1.1.2 send-community
router bgp 1 address-family l2vpn-evpn neighbor 2.2.2.2 send-community
router bgp 1 address-family l2vpn-evpn neighbor 1.1.1.2 send-community ext
router bgp 1 address-family l2vpn-evpn neighbor 2.2.2.2 send-community ext
router bgp 1 address-family l2vpn-evpn neighbor 1.1.1.2 next-hop-unchanged
router bgp 1 address-family l2vpn-evpn neighbor 2.2.2.2 next-hop-unchanged
router bgp 1 redistribute connected
router bgp 1 address-family l2vpn-evpn vni auto-create
```

2. Configure the MLAG Standby in the following way:

```
ip routing vrf default
protocol nve
interface nve 1
interface nve 1 vxlan source interface loopback 1
interface nve 1 nve controller bgp
interface nve 1 vxlan mlag-tunnel-ip 5.5.5.5
interface nve 1 nve vni auto-vlan-map base 100000
```

```
interface mlag-port-channel 1 switchport access vlan 10
interface vlan 10 ip address 10.0.0.252/24 primary
interface vlan 4094 ip address 2.2.2.2/24 primary

protocol magp
interface vlan 10 magp 10
interface vlan 10 magp 10 ip virtual-router address 10.0.0.1
interface vlan 10 magp 10 ip virtual-router mac-address 00:00:5e:00:00:10

mlag-vip MLAG-1 ip 11.11.11.11 /24 force
interface port-channel 1 ipl 1
interface vlan 4094 ipl 1 peer-address 2.2.2.1

protocol bgp
router bgp 1 vrf default
router bgp 1 router-id 200.0.1.1 force
router bgp 1 neighbor 1.1.1.2 remote-as 1
router bgp 1 neighbor 2.2.2.1 remote-as 1
router bgp 1 address-family l2vpn-evpn neighbor 1.1.1.2 send-community
router bgp 1 address-family l2vpn-evpn neighbor 2.2.2.1 send-community
router bgp 1 address-family l2vpn-evpn neighbor 1.1.1.2 send-community ext
router bgp 1 address-family l2vpn-evpn neighbor 2.2.2.1 send-community ext
router bgp 1 address-family l2vpn-evpn neighbor 1.1.1.2 next-hop-unchanged
router bgp 1 address-family l2vpn-evpn neighbor 2.2.2.1 next-hop-unchanged
router bgp 1 redistribute connected
router bgp 1 address-family l2vpn-evpn vni auto-create
```

14.7 EVPN Logging Examples

14.7.1 EVPN MAC Mobility Logs

MAC mobility warning is detected when a MAC address is noticed to move between a local and one or more remote customer site 5 times in a period of 180 seconds. This indicates that multiple hosts have been configured with the same MAC address. The MAC mobility warning is cleared when only one route for the MAC address is left (either local or remote).

When detecting EVPN MAC duplication, the following message will appear:

```
[metad.WARNING]: EVPN MAC duplication detected for MAC 24:8A:07:A0:B0:0D, IP 2.2.2.2 and VLAN 6 from BGP neighbor
1.1.1.1
```

A static MAC error is detected when a remote route is received for a MAC address for which a local existing route has been marked as static. The local route being marked as static indicates that the MAC address is not expected to move. In this case, any remote route with this MAC address is an error. The static MAC error is cleared when all remote routes for the MAC address are withdrawn or if the local route is no longer marked as static.

When receiving EVPN MAC mobility route for a static MAC address, the following message will appear:

```
[metad.WARNING]: EVPN MAC mobility route received for sticky MAC 24:8A:07:A0:B0:0D, IP 2.2.2.2 and VLAN 6 from BGP
neighbor 1.1.1.1
```

15 IP Routing

The following pages provide information on configuring IP routing (L3) protocols and features.

- IP Routing Overview
- OSPF
- BGP
- Bidirectional Forwarding Detection (BFD) Infrastructure
- Policy Rules
- VRRP
- MAGP
- DHCP Relay

15.1 IP Routing Overview

15.1.1 IP Interfaces

NVIDIA Onyx supports the following 3 types of IP interfaces:

- VLAN interface
- Loopback interface
- Router port interface

Onyx supports up to 999 IP interfaces.

Each IP interface can be configured with multiple IP addresses. The first address assigned to the interface automatically becomes its primary address (only one primary address is supported per interface), and the rest are secondary addresses.

> Secondary addresses are advertised via OSPF. No “HELLO” messages are sent on them and no adjacencies are established on them either.

Primary addresses cannot be modified once assigned. To assign a different primary address, all addresses of the interface must be removed and then reconfigured.

Up to 16 IPv4 (as well as IPv6) addresses are supported on each IP interface.

> IPv4 link local IP addresses such as 169.254.x.x can be assigned to IP interfaces, thus allowing all routing, forwarding functions and applications on top of the interfaces to function as the real IP addresses. Only unique addresses from that range can be assigned to IP interface, same address assignment is not supported.

Since 169.254.101.101 is already used as BGP unnumbered neighbor address, it is recommended not to use this address in the network if BGP unnumbered neighbor is to ever be enabled.

15.1.1.1 VLAN Interfaces

VLAN interface is a logical IPv4 interface created per subnet over a specific 802.1Q VLAN ID. If two hosts from two different subnets need to communicate (via the IP layer), the network administrator needs to configure two interface VLANs, one for each of the subnets.

Each interface VLAN has the following attributes:

- Admin state
- Operational state
- MAC address
- IP address and mask
- MTU
- Description
- Set of counters

15.1.1.2 Loopback Interfaces

Loopback interface is a logical software entity where traffic transmitted to this interface is immediately received on the sending end.

15.1.1.3 Router Port Interfaces

Router port interface is a regular switch port configured to operate as an L3 interface. Router port interfaces are assigned an IP address and all L3 commands become applicable to them.

Once configured, router port interfaces no longer partake in the bridging activities of the switch and VLANs configured on them are separate from the pool allocated for the switch ports.

15.1.1.4 Configuring a VLAN Interface

1. Create a VLAN. Run:

```
switch (config)# vlan 10
switch (config vlan 10)# exit
```

2. Assign a physical interface to this VLAN. Run:

```
switch (config)# interface ethernet 1/1
switch (config interface ethernet 1/1)# switchport mode access
switch (config interface ethernet 1/1)# exit
```

3. There must be at least one interface in the operational state “UP”. Run:

```
switch (config)# show interface ethernet 1/1 status
Port                 Operational state          Speed               Negotiation
----                 -----------------          -----               -----------
Eth1/1               Up                         40 Gbps             No-Negotiation
```

4. Create a VLAN interface that matches the VLAN. Run:

```
switch (config)# interface vlan 10
switch (config interface vlan 10)#
```

5. Configure an IP address and a network mask to the interface. Run:

```
switch (config interface vlan 10)# ip address 10.10.10.10 /24
```

6. Verify VLAN interface configuration. Run:

```
switch (config interface vlan 10) # show interfaces vlan 10

Vlan 10:
  Admin state      : Enabled
  Operational state: Down
  Autostate        : Enabled
  Mac Address      : 24:8a:07:f3:04:c8
  DHCP client      : Disabled

  IPv4 address:
    10.10.10.10/24 [primary]

  Broadcast address:
    10.10.10.255 [primary]

  Arp responder: Disabled
  MTU          : 1500 bytes
  Arp timeout  : 1500 seconds
  Icmp redirect: Enabled
  Description  : my-ip-interface
  VRF          : default
  Counters     : Disabled
```

15.1.1.5 Configuring a Loopback Interface

1. Create a loopback interface. Run:

```
switch (config)# interface loopback 2
switch (config interface loopback 2)#
```

2. Configure an IP address on the loopback interface. Run:

```
switch (config interface loopback 2)# ip address 20.20.20.20 /32
```

3. Verify loopback interface configuration. Run:

```
switch (config interface loopback 2)# show interfaces loopback 2

Loopback 2:
  IPv4 address:
    20.20.20.20/32 [primary]

  Broadcast address:
    20.20.20.20 [primary]

  MTU        : 1500 bytes
  Description: my-loopback
  VRF        : default
```

15.1.1.6 Configuring a Router Port Interface

1. Enter an Ethernet interface's configuration context. Run:

```
switch (config)# interface ethernet 1/10
switch (config interface ethernet 1/10)#
```

2. Configure the Ethernet interface to become an router port interface. Run:

```
switch (config interface ethernet 1/10)# no switchport force
```

3. Configure an IP address on the router port interface. Run:

```
switch (config interface ethernet 1/10)# ip address 100.100.100.100 /24
```

4. Verify router port interface configuration. Run:

```
switch (config interface ethernet 1/10)# show interfaces ethernet 1/10

Eth1/10:
  Admin state                     : Enabled
  Operational state               : Down
  Last change in operational status: Never
  Boot delay time                 : 0 sec
  Description                     : N/A
  Mac address                     : 24:8a:07:f3:04:c8
  MTU                             : 1500 bytes (Maximum packet size 1522 bytes)
  Fec                             : auto
  Flow-control                    : receive off send off
  Supported speeds                : 1G 10G 25G
  Advertised speeds               : 1G 10G 25G
  Actual speed                    : Unknown
  Auto-negotiation                : Enabled
  Width reduction mode            : Unknown
  DHCP client                     : Disabled
  Autoconfig                      : Disabled

  IPv4 address:
    100.100.100.100/24 [primary]

  Broadcast address:
    100.100.100.255 [primary]

  Arp responder: Disabled
  Arp timeout    : 1500 seconds
  VRF            : default
  Forwarding mode: inherited cut-through

  Telemetry sampling: Disabled   TCs: N\A
  Telemetry threshold: Disabled  TCs: N\A
  Telemetry threshold level: N\A

  Last clearing of "show interface" counters: Never
  60 seconds ingress rate                   : 0 bits/sec, 0 bytes/sec, 0 packets/sec
  60 seconds egress rate                    : 0 bits/sec, 0 bytes/sec, 0 packets/sec


  Rx:
    0                     packets
    0                     unicast packets
    0                     multicast packets
    0                     broadcast packets
    0                     bytes
    0                     discard packets
    0                     error packets
    0                     fcs errors
    0                     undersize packets
    0                     oversize packets
    0                     pause packets
    0                     unknown control opcode
    0                     symbol errors

  Tx:
    0                     packets
    0                     unicast packets
    0                     multicast packets
    0                     broadcast packets
    0                     bytes
    0                     discard packets
    0                     error packets
    0                     hoq discard packets
```

15.1.2 Equal Cost Multi-Path Routing (ECMP)

Equal-cost multi-path routing (ECMP) is a routing strategy where next-hop packet forwarding to a single destination can occur over multiple paths.

In the following figures, routers R1 and R2 can both access each of their router peer networks. Router R1 routing table for 10.0.40/24 will contain the following routes:

- 10.0.10.2
- 10.0.20.2
- 10.0.30.2

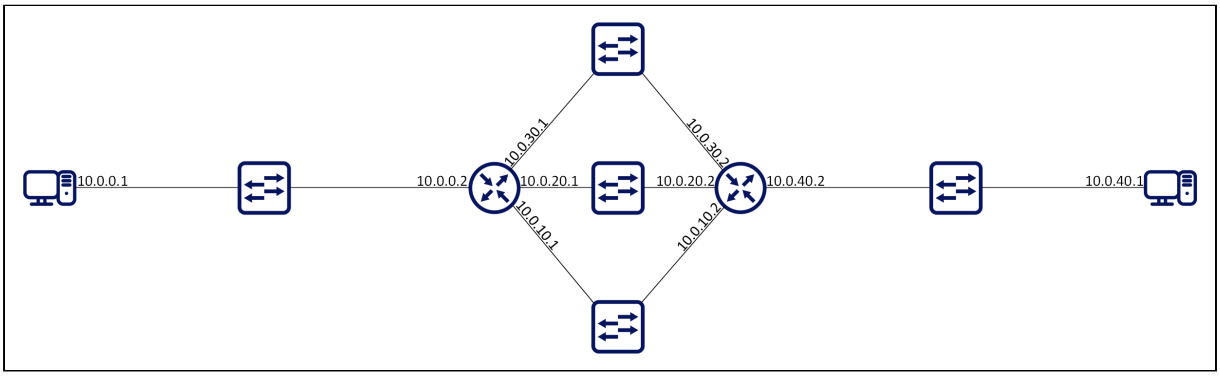

The load balancing function of the ECMP is configured globally on the system.

Hash algorithm can be symmetric or asymmetric. In symmetric hash functions bidirectional flows between routes will follow the same path, while in asymmetric hash functions, bidirectional traffic can follow different paths in both directions.

The following load balancing types are supported:

- Source IP & Port – source IP (SIP) and source UDP/TCP port: If the packet is not UDP/TCP, only SIP is used for the hash calculation. This is an asymmetric hash function.
- Destination IP & Port – destination IP (DIP) and destination UDP/TCP port: If the packet is not UDP/TCP, only DIP is used for the hash calculation. This is an asymmetric hash function.
- Source and Destination IP & Port – destination and source IP, as well as destination and source UDP/TCP port: If the packet is not UDP/TCP, only SIP/DIP are used for the hash calculation. This is a symmetric hash function.
- Traffic Class – Load balance based on the traffic class assigned to the packet. This is an asymmetric hash function.
- All (default) – all above fields are part of the hash calculations. This is a symmetric hash function.

15.1.2.1 Hash Functions

It is advised that LAG and ECMP hash function configuration over more than one hop is different. If the same hash function is used over two hops, all the traffic sorted from one hop to following one will arrive already having the same characteristics, which will render the next hash function useless. For example, configure load-balancing on the first hop based on source IP while on the next hop based on destination IP.

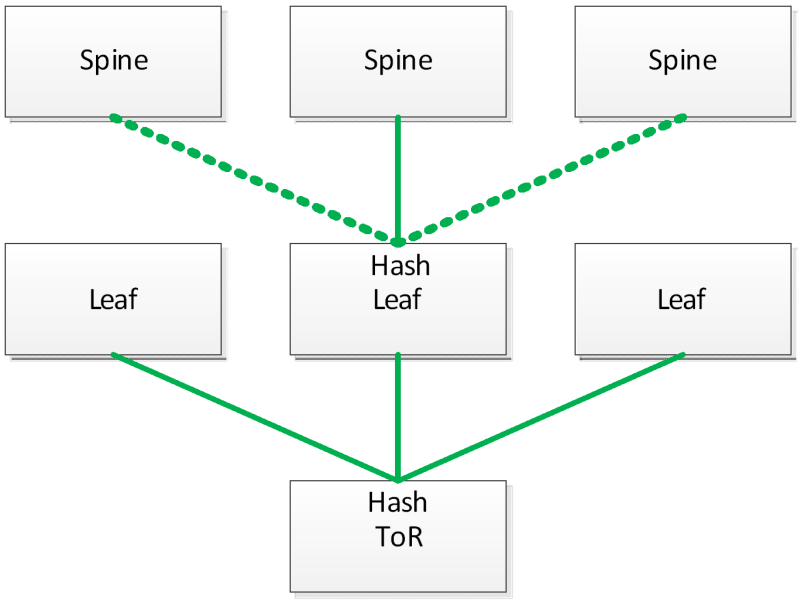

15.1.2.2 ECMP Consistent Hashing

In an IP network multiple flows share the same path defined by their destination prefix. ECMP allows those flows to travel with the same prefix and be distributed over multiple next hops that usually belong to different physical links, in order to reach better bandwidth utilization. When using the standard ECMP some links in the network become unreachable, thus the next hop list and hash function distribution change, and flows are moved to other links. Packet reordering in the network or failure in a user session might occur, while others which use anycast IP addresses utilize ECMP distribution for load balancing. Therefore, changing the next hop may cause flows to arrive to the wrong destination.

When network is reconfigured, and route next hop set is changed, flows that are not affected by the change should continue to be sent to the same next hops and keep the same outgoing link.

Using consistent hash containers enables you to use size arrays with next hop buckets to make sure unaffected flows are sent to the same next hops when some next hops are removed from the container. When a new next hop is added to the consistent hash container, some buckets are replaced with a new next hop, so part of the existing flows are moved to a new next hop.

When a route is installed, it points to a hash container. Each flow in the route is mapped to a respective bucket, and is eventually forwarded to the next hop in the bucket.

In the following example we see a single route with 3 flows and 4 next hops, so the container has 12 bucket.

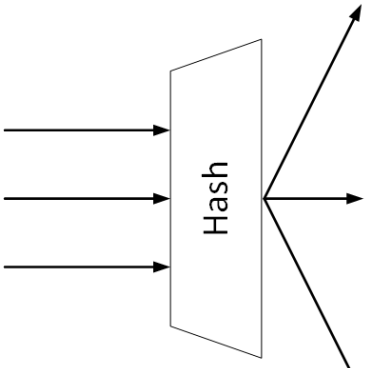

0	Next Hop 1
1	Next Hop 2
2	Next Hop 3
3	Next Hop 4
4	Next Hop 1
5	Next Hop 2
6	Next Hop 3
7	Next Hop 4
8	Next Hop 1
9	Next Hop 2
10	Next Hop 3
11	Next Hop 4

15.1.2.2.1 Remove Next Hops

Unlike the default IP load-sharing hashing, when consistent hashing is used, and a next hop needs to be removed, the number of hash buckets does not change. All appearances of the deleted next hop are removed from the container and replaced by the remaining next hops.

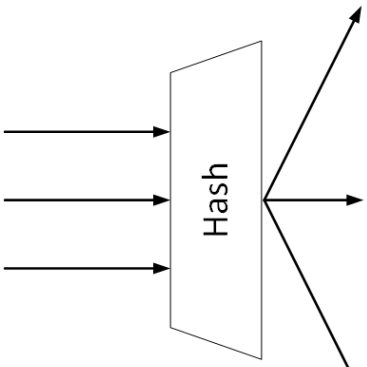

0	Next Hop 1
1	~~Next Hop 2~~ Next Hop 1
2	Next Hop 3
3	Next Hop 4
4	Next Hop 1
5	~~Next Hop 2~~ Next Hop 3
6	Next Hop 3
7	Next Hop 4
8	Next Hop 1
9	~~Next Hop 2~~ Next Hop 4
10	Next Hop 3
11	Next Hop 4

15.1.2.2.2 Add Next Hops

When adding a new next hop, some of existing next hops should be removed from the hash, and the new next hop should be located in one of the newly available places.The new next hops are not applied to HW immediately, but only after a convergence time period.

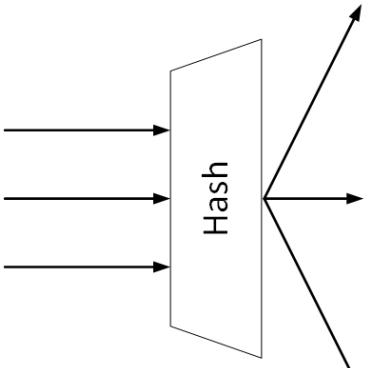

0	~~Next Hop 1~~ Next Hop 5
1	Next Hop 2
2	Next Hop 3
3	Next Hop 4
4	Next Hop 1
5	~~Next Hop 2~~ Next Hop 5
6	Next Hop 3
7	Next Hop 4
8	Next Hop 1
9	Next Hop 2
10	~~Next Hop 3~~ Next Hop 5
11	Next Hop 4

15.1.2.2.3 Supported Number of Containers

When the consistent hashing containers count exceeds the maximum number of containers, the operational state of consistent hashing function will become “unstable” and the containers with the same next hop sets will be merged to release more resources. Once more resources are available to deploy the containers, the operation state will become “stable”.

In the unstable case which may result from lack of consistent hashing resources, the new route will be installed as a non-consistent route, and a random next hop from its next hop set will be chosen as the actual next hop and installed in hardware. The route will only be partially programed in hardware.

Container Bucket Size	Default Number of Containers	Maximum Number of Containers
512	40	96
1024	20	48

15.1.2.2.4 Configuring Consistent Hashing

To configure consistent hashing, run “ip load-sharing type consistent”.

15.1.2.3 Virtual Routing and Forwarding

Virtual Routing and Forwarding (VRF) allows multiple routing table instances to coexist within the same router simultaneously. Since the routing instances are independent, IP addresses on each routing table may overlap without conflicting with each other.

VRF can be used for the following purposes:

- Ensure customer privacy and security
- Separate between management and user data
- Support customers with the same address space
- Support VPN

Multiple routing instances defined in the router can have different purposes and can be configured in different manners:

- Different IP interfaces can be attached to different VRFs (only one IP interface can be in a single VRF)
- Routing in VRF can be enabled or disabled
- Each VRF component can run its own routing protocol independently from other instances
- Differently configured IPv4 and IPv6 services

The first VRF in the system is created automatically and it is called “default” VRF. It cannot be deleted or configured.

Onyx supports up to 64 VRFs, 8 instances of BGP, and 8 instances of OSPF.

15.1.3 ARP Neighbor Discovery Responder

ARP functionality in IP/Ethernet networks is needed to provide mapping from IP addresses to L2 MAC addresses. This request may be sent in multiple cases:

- A station wants to initiate an IP session with another station on the same IP subnet and needs to obtain its L2 address
- A station wants to update other stations that its MAC address has changed
- A station wants to check that the MAC address of its peer did not change
- The peer responds with unicast ARP response.

The following are two scenarios when ARP responder functionality is needed:

- Network wants to avoid broadcast in the network or on some parts of the network, so broadcast ARP packets are not distributed in that part of the network
- There is no L2 connectivity between some parts of the network, and even IP addressing scheme does not reflect it

ARP responder answers a broadcast ARP requests that arrive to the switch.

ARP responder is configured on an IP interface (with or without IP address) of any type (e.g. VLAN interface, router port, or LAG).

> Only IP interfaces in UP admin state respond to ARP.

This functionality is provided for all ARP entries that are configured or provided on the interface: Static, dynamic, or per protocol.

> There is no need to enable IP routing in the system to enable ARP responder functionality.

If a user has multiple VRFs the interface can be created in any VRF. If IP routing is disabled the interface is created in default VRF.

ARP responder can be enabled together with IP routing and given an interface which can be used in routing.

When IP routing on the interface is enabled, all entries that have been used by the responder become ARP entries for the router and vice versa.

> A user must avoid using ARP responder in broadcast networks—the system itself does not block it.

15.1.3.1 Configuring ARP Responder

In order to initialize ARP responder:

1. Create IP interface. Run:

```
switch (config) # interface vlan 10
switch (config interface vlan 10) #
```

2. Initialize ARP responder on the interface. Run:

```
switch (config interface vlan 10) # ip arp responder
```

3. Create static ARP entries on VLAN. Run:

```
switch (config interface vlan 10) # ip arp 172.130.11.1 00:11:22:33:44:55
```

4. Create ACL to drop broadcast, and assign it to all relevant L2 interface (VLAN's members). Run:

```
switch (config interface vlan 10) # mac access-list new
switch (config interface vlan 10) # mac access-list new seq-number 10 deny any FF:FF:FF:FF:FF:FF mask
FF:FF:FF:FF:FF:FF
switch (config interface vlan 10) # interface ethernet 1/3-1/5 mac port access-group new
```

15.1.4 Policy Based Routing (PBR)

Usually layer 3 forwarding is done based on destination IP: a router will extract packet destination IP from the packet header, match it to its routing table in Longest prefix match order, and forward it according the lookup result. In some cases, it is required that the routing decision will depend on different criteria such as source IP, source or destination port, packet type, and so forth.

PBR provides a way to implement such behavior. PBR is implemented as match/action table and influence the destination to which a packet should go based on various packet fields and not only based on the destination IP address.

PBR is applied to ingress packets after Ingress ACL and OpenFlow rules for packets that are eligible for routing.

15.1.5 General IP Routing Commands

- 15.1.5.1 vrf definition
- 15.1.5.2 routing-context vrf
- 15.1.5.3 ip routing
- 15.1.5.4 description
- 15.1.5.5 rd
- 15.1.5.6 vrf forwarding
- 15.1.5.7 clear ip routing counters
- 15.1.5.8 show ip routing
- 15.1.5.9 show ip routing counters
- 15.1.5.10 show routing-context vrf
- 15.1.5.11 show vrf
- 15.1.5.12 IP Interface
  - 15.1.5.12.1 switchport
  - 15.1.5.12.2 encapsulation dot1q vlan
  - 15.1.5.12.3 interface ip enable
- 15.1.5.13 Interface VLAN
  - 15.1.5.13.1 interface vlan
  - 15.1.5.13.2 interface vlan no-autostate
  - 15.1.5.13.3 ip address
  - 15.1.5.13.4 counters
  - 15.1.5.13.5 description
  - 15.1.5.13.6 mtu
  - 15.1.5.13.7 shutdown
  - 15.1.5.13.8 clear counters
  - 15.1.5.13.9 ip icmp redirect
  - 15.1.5.13.10 show interfaces
  - 15.1.5.13.11 show interfaces vlan
  - 15.1.5.13.12 show ip interface
  - 15.1.5.13.13 show ip interface brief
  - 15.1.5.13.14 show interfaces configured
  - 15.1.5.13.15 show ip
  - 15.1.5.13.16 show ip interface mgmt0
  - 15.1.5.13.17 show ip interface port-channel
  - 15.1.5.13.18 show ip interface vrf
  - 15.1.5.13.19 show ip interface vrf vrf
  - 15.1.5.13.20 show ipv6 interface
  - 15.1.5.13.21 show ipv6 interface brief
  - 15.1.5.13.22 show ipv6
  - 15.1.5.13.23 show ipv6 interface loopback
  - 15.1.5.13.24 show ipv6 interface port-channel
  - 15.1.5.13.25 show ipv6 interface vlan

        - 15.1.5.13.26 show ipv6 interface vrf
        - 15.1.5.13.27 show ipv6 interface vrf brief
    - 15.1.5.14 Loopback Interface
        - 15.1.5.14.1 interface loopback
        - 15.1.5.14.2 interface vrf ip address alias
        - 15.1.5.14.3 ip address
        - 15.1.5.14.4 description
        - 15.1.5.14.5 show interfaces loopback
    - 15.1.5.15 Routing and ECMP
        - 15.1.5.15.1 ip route
        - 15.1.5.15.2 ip load-sharing
        - 15.1.5.15.3 show ip route
        - 15.1.5.15.4 show ip route vrf
        - 15.1.5.15.5 show ip route -a
        - 15.1.5.15.6 show ip route failed
        - 15.1.5.15.7 show ip route static
        - 15.1.5.15.8 show ip route static multicast-override
        - 15.1.5.15.9 show ip route summary
        - 15.1.5.15.10 show ip route interface
        - 15.1.5.15.11 show ip load-sharing
    - 15.1.5.16 Network to Media Resolution (ARP)
        - 15.1.5.16.1 ip arp
        - 15.1.5.16.2 ip arp responder
        - 15.1.5.16.3 ip arp timeout
        - 15.1.5.16.4 clear ip arp
        - 15.1.5.16.5 show ip arp
    - 15.1.5.17 IP Diagnostic Tools
        - 15.1.5.17.1 ping
        - 15.1.5.17.2 traceroute
        - 15.1.5.17.3 tcpdump
    - 15.1.5.18 QoS
        - 15.1.5.18.1 qos map dscp-to-pcp preserve-pcp
    - 15.1.5.19 PBR
        - 15.1.5.19.1 nexthop-group direct
        - 15.1.5.19.2 nexthop-group direct nexthop interface
        - 15.1.5.19.3 nexthop-group recursive nexthop
        - 15.1.5.19.4 route-map
        - 15.1.5.19.5 route-map sequence match rule
        - 15.1.5.19.6 route-map sequence nexthop-group
        - 15.1.5.19.7 route-map sequence counter
        - 15.1.5.19.8 bind/unbind route-map on interface
        - 15.1.5.19.9 show nexthop-groups
        - 15.1.5.19.10 show route-maps
        - 15.1.5.19.11 route-map to interface bind
        - 15.1.5.19.12 show pbr general information

15.1.5.1 vrf definition

	vrf definition <vrf-name> [force] no vrf definition <vrf-name> [force] Creates the VRF.	
Syntax Description	vrf-name	VRF session name
	force	"force" option was added on VRF creation command to bypass user confirmation for creating "mgmt" VRF
Default	N/A	
Configuration Mode	config	
History	3.4.2008	
	3.6.6000	Updated the notes section
	3.9.2000	Added force option
Example	`switch (config) # vrf definition my-vrf` `switch (config vrf definition my-vrf) #`	
Related Commands		
Notes	• 63 VRFs are supported aside from the default VRF • In case of "mgmt" VRF creation, CLI will ask permission to save the current configuration (behave like "configuration write") and reload the system. If "force" option was passed, then no confirmation is needed. After reboot, mgmt VRF will be created with management interfaces in it. Also, clustered, mDNS, OpenFlow API and FTP/TELNET servers will run in managemtn VRF when started. All other services does not change their existing configuration. In case of management VRF removal, the CLI will ask permission to remove services that running in management VRF, save new configuration, and reboot the switch. If "force" option was passed, no confirmation is needed. After reboot, mgmt VRF will be removed and management interfaces will be moved to "default" VRF. Also, clusterd, mDNS, OpenFlow API, and FTP/TELNET servers will run in "default" VRF when started. Other services that were enabled in management VRF will be disabled, except ones that are enabled by default (i.e., "ntp", "snmp-server", "tacacs-server", "radius-server", "ldap", "web", and so forth)—they will be reset and enabled in "default" VRF. The logic of moving/shutting down services from removed VRF could be applied for ALL user-defined VRF`s.	

15.1.5.2 routing-context vrf

	routing-context vrf <vrf-name> Enters the active-context of the specified session.	
Syntax Description	vrf-name	VRF session name
Default	N/A	
Configuration Mode	config	
History	3.4.2008	
Example	`switch (config) # routing-context vrf my-vrf`	

Related Commands	
Notes	• If a routing-context is configured, the user does not have to explicitly specify the VRF name parameter in this or any other VRF command • If no routing-context is configured and the user does not specify the VRF name, default VRF is used

15.1.5.3 ip routing

	ip routing [vrf <vrf-name>] Enables L3 forwarding between high speed interfaces.	
Syntax Description	vrf-name	VRF session name
Default	N/A	
Configuration Mode	config	
History	3.4.1802	
	3.4.2008	Added VRF parameter
Example	`switch (config) # ip routing vrf my-vrf`	
Related Commands		
Notes	• RD must be configured to enable IP routing on the VRF • If no routing-context is specified, the “routing-context” VRF is automatically configured.	

15.1.5.4 description

	description <description> no description forceAdds description for the VRF. The no form of the command removes the description of the VRF.	
Syntax Description	description	Text string
	force	Forces deletion (no confirmation needed if configuration exists inside the VRF)
Default	N/A	
Configuration Mode	config vrf definition	
History	3.4.2008	
Example	`switch (config vrf definition my-vrf) # description vrf-description`	
Related Commands		
Notes		

15.1.5.5 rd

	rd [<ip addr>:<0-65,535> | <AS Number>:<0-4,294,967,295> | <AS Number>:<ip addr>] Adds a Route Distinguisher (RD) to the VRF configuration mode.	
Syntax Description	ip-addr	IPv4 address

	AS Number	Asynchronous machine number
Default	N/A	
Configuration Mode	config vrf definition	
History	3.4.2008	
Example	`switch (config vrf definition my-vrf) # rd 10.10.10.10:2`	
Related Commands		
Notes	• RDs internally identify routes belonging to a VRF to distinguish overlapping or duplicate IP address ranges. This allows the creation of distinct routes to the same IP address for different VPNs. The RD is a 64-bit number made up of an AS number or IPv4 address followed by a user-selected ID number. Once an RD has been assigned to a VRF it cannot be changed. To change the RD, remove the VRF then create it again. VRF is not active until an RD is defined. • An RD must be defined to enable IP routing on the VRF	

15.1.5.6 vrf forwarding

	vrf forwarding <vrf-name> Maps an interface to VRF.	
Syntax Description	vrf-name	VRF session name
Default	N/A	
Configuration Mode	config interface ethernet set as router port interface config interface vlan config interface loopback	
History	3.4.2008	
Example	`switch (config interface ethernet 1/2) # vrf forwarding my-vrf`	
Related Commands		
Notes		

15.1.5.7 clear ip routing counters

	clear ip routing counters Clears counters, related to NULL interface and dropped packets by router.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6102
Example	`switch (config) # clear ip routing counters`
Related Commands	
Notes	

15.1.5.8 show ip routing

	show ip routing [vrf <vrf-name> | all] Displays IP routing information per VRF.	
Syntax Description	vrf	Displays information for specific VRF
	all	Displays information on all VRFs
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.0230	
	3.4.2008	Added VRF parameter
	3.6.8008	Updated example
	3.9.0500	Updated example

Example	`switch (config) # show ip routing`
	`VRF Name default:`
	`IP routing : disabled`
	`Global virtual router mac: aa:bb:cc:00:00:11`
	
	`switch (config) # show ip routing vrf all`
	`VRF Name default:`
	`IP routing: enabled`
	`VRF Name new:`
	`IP routing: disabled`	

Related Commands		
Notes	If no routing-context is specified, the “routing-context” VRF is automatically displayed.	

15.1.5.9 show ip routing counters

	show ip routing [vrf <vrf-name> | all] counters Display counters, related to NULL interface and dropped packets by router.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6102	
	3.9.1300	Updated example

Example	`switch (config) # show ip routing counters`
	
	`0 packets discarded by router`
	`0 bytes discarded by router`
	`0 packets to null interface`
	`0 bytes to null interface`	

Related Commands		
Notes		

15.1.5.10 show routing-context vrf

	show routing-context vrf Displays VRF active context.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.4.2008
Example	`switch (config) # show routing-context vrf` `VRF active context: my-vrf`
Related Commands	
Notes	

15.1.5.11 show vrf

	show vrf [<vrf-name> | all] Displays VRF information.	
Syntax Description	all	Displays information for all VRF instances
	vrf-name	Name of VRF instance
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.6.6000	Updated example
	3.9.1900	Updated example

Example	`switch (config) # show vrf my-vrf`
	`VRF Info:`
	
	`Name: default`
	`RD: NA`
	`Description: NA`
	`IP routing state: Enabled`
	`IPv6 routing state: Disabled`
	`IP multicast routing state: Enabled`
	`Protocols: IPv4, PIM-SM`
	`Interfaces: Eth1/1`
	
	`switch (config) # show vrf my-vrf`
	`VRF Info:`
	
	`Name: default`
	`RD: NA`
	`Description: NA`
	`IP routing state: Enabled`
	`IPv6 routing state: Disabled`
	`IP multicast routing state: Enabled`
	`Protocols: IPv4, PIM-BIDIR`
	`Interfaces: Eth1/1`	

Related Commands		
Notes	If no routing-context is specified, the “routing-context” VRF is automatically displayed.	

15.1.5.12 IP Interface

15.1.5.12.1 switchport

	switchport [force] no switchport [force] Configures the Ethernet interface as a regular switchport. The no form of the command configures the Ethernet interface as router port interface.	
Syntax Description	force	Forces configuration even if the interface's admin state is enabled
Default	N/A	
Configuration Mode	config interface ethernet config interface port-channel	
History	3.3.5200	
	3.6.4006	Added storm-control support

Example	`switch (config interface ethernet 1/10)# no switchport force`
	
	`error message is case storm-control is configured on port:`
	
	`% interface * has storm control configuration. Please remove it first`	

Related Commands		
Notes	• When storm-control is configured on port, an error message will appear • Force command deletes all storm-control configuration from port	

15.1.5.12.2 encapsulation dot1q vlan

	encapsulation dot1q vlan <vlan-id> [force] no encapsulation dot1q vlan [force] Enables L2 802.1Q encapsulation of traffic on a specified router port interface in a VLAN. The no form of the command disables L2 802.1Q encapsulation of traffic on a specified router port interface in a VLAN.	
Syntax Description	vlan-id	Enables L2 802.1Q encapsulation of traffic on a router port interface in a VLAN
	force	Forces admin state down
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.3.5200	
Example	`switch (config interface ethernet 1/10)# encapsulation dot1q vlan 10`	
Related Commands		
Notes		

15.1.5.12.3 interface ip enable

	interface <vlan | ethernet | port-channel> <ifname> ip enable no interface <vlan | ethernet | port-channel> <ifname> ip enable Enables IP forwarding on the interface. The no form of the command disables IP forwarding on the interface.	
Syntax Description	vlan	VLAN type interface
	ethernet	Ethernet type interface
	port-channel	LAG type interface
	ifname	interface id
Default	Disabled	
Configuration Mode	config	
History	3.9.0300	
Example	`switch (config interface vlan 10)# ip enable`	
Related Commands	show ip interface vrf	
Notes		

15.1.5.13 Interface VLAN

15.1.5.13.1 interface vlan

	interface vlan <vid> no interface vlan <vid> Creates a VLAN interface and enters the interface VLAN configuration mode. The no form of the command deletes the VLAN interface.	
Syntax Description	vid	VLAN ID
Default	N/A	
Configuration Mode	config	
History	3.2.0230	
Example	`switch (config) # interface vlan 10` `switch (config interface vlan 10) #`	
Related Commands	ip routing vlan <vlan-id> switchport mode switchport access show interface vlan	
Notes	• Make sure the VLAN was created, using the command “vlan <vlan-id>” in the global configuration mode • The VLAN must be assigned to one of the L2 interfaces. To do so, run the command “swichport ...” • At least one interface belong to that VLAN must be in UP state	

15.1.5.13.2 interface vlan no-autostate

	interface vlan <vid> no-autostate no interface vlan <id> no-autostate Disables the VLAN interface autostate such that its operational state remains up as long as its admin state is up, even if no port in the relevant VLAN egress-list is operationally up. The no form of the command enables this functionality.	
Syntax Description	vid	
Default	N/A	
Configuration Mode	config	
History	3.6.4006	
Example	`switch (config) # interface vlan 10 no-autostate` `switch (config) # interface vlan 10-13 no-autostate`	
Related Commands	show ip interface vlan	
Notes		

15.1.5.13.3 ip address

	ip address <ip-address> <mask> no ip address [<ip-address> [<mask>]] Enters user-defined IPv4 address for the interface. The no form of the command removes the specified IPv4 address. If no address is specified, then all IPv4 addresses of this interface are removed.	
Syntax Description	ip-address	
	mask	There are two possible ways to the mask: • /length (i.e. /24) • Network address (i.e. 255.255.255.0) The mask length may be configured without a space (i.e. <ipv4-address>/<length>)
Default	0.0.0.0/0	
Configuration Mode	config interface vlan	
History	3.2.0230	
Example	`switch (config interface vlan 10) # ip address 10.10.10.10 /24`	
Related Commands	interface vlan show interfaces vlan	
Notes	An interface may have up to 16 IPv4 address assignments	

15.1.5.13.4 counters

	counters no counters Enables counters on the IP interface. The no form of the command disables counters gathering on the IP interface.
Syntax Description	N/A
Default	Disabled

Configuration Mode	config interface vlan
History	3.2.0230
Example	`switch (config interface vlan 10) # counters`
Related Commands	interface vlan show interfaces vlan
Notes	• Enabling counters for the router interface adds delay to the traffic stream • There are maximum of 16 counter sets

15.1.5.13.5 description

	description <string> no description Enters a description for the interface. The no form of the command sets the description to default.	
Syntax Description	string	User defined string
Default	“”	
Configuration Mode	config interface vlan	
History	3.2.0230	
Example	`switch (config interface vlan 10) # description my-ip-interface`	
Related Commands	interface vlan show interfaces vlan	
Notes		

15.1.5.13.6 mtu

	mtu <size> [force] no mtu Sets the Maximum Transmission Unit for the interface. The no form of the command sets the MTU to default.	
Syntax Description	size	Range: 1500-9216 bytes
Default	9216 bytes	
Configuration Mode	config interface vlan	
History	3.2.0230	
	3.9.2000	Changed default MTU size and added note
Example	`switch (config interface vlan 10)# mtu 9216`	
Related Commands	interface vlan show interfaces vlan	
Notes	In switches that perform upgrade to version 3.9.2000, existing L3 interfaces will stay with MTU 1500 (or any other value that was configured). Newly created interfaces will be created with MTU 9216 (the new default).	

15.1.5.13.7 shutdown

	shutdown no shutdown Disables the interface. The no form of the command enables the interface.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config interface vlan
History	3.1.0000
Example	`switch (config interface vlan 20) # shutdown`
Related Commands	interface vlan show interfaces vlan
Notes	

15.1.5.13.8 clear counters

	clear counters Clears the interface counters.
Syntax Description	N/A
Default	N/A
Configuration Mode	config interface vlan
History	3.2.0230
Example	`switch (config interface vlan 10) # clear counters`
Related Commands	counters interface vlan show interfaces vlan
Notes	

15.1.5.13.9 ip icmp redirect

	ip icmp redirect no ip icmp redirect Enables ICMP redirect. The no form of the command disables ICMP redirect.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config interface vlan
History	3.4.0010
Example	`switch (config interface vlan 10) # no ip icmp redirect`
Related Commands	interface vlan show interfaces vlan

Notes	ICMP redirect transmits messages to hosts alerting them about the existence of more efficient routes to a specific destination

15.1.5.13.10 show interfaces

	show interfaces [brief] Displays interface configuration.	
Syntax Description	brief	Displays brief output
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.3000	
	3.6.8008	Updated example

Example

```
switch (config) # show interfaces

Interface lo status:
  Comment          :
  Admin up         : yes
  Link up          : yes
  DHCP running     : no
...
Interface mgmt0 status:
  Comment          :
  Admin up         : yes
  Link up          : yes
  DHCP running     : yes
...
Interface mgmt1 status:
  Comment          :
  Admin up         : yes
  Link up          : yes
  DHCP running     : yes (but no valid lease)
...
Eth1/1:
  Admin state                       : Enabled
  Operational state                 : Up
  Last change in operational status: 0:22:11 ago (5 oper change)
  Boot delay time                   : 0 sec
...
```

Related Commands	interface vlan show interfaces vlan
Notes	ICMP redirect transmits messages to hosts alerting them about the existence of more efficient routes to a specific destination

15.1.5.13.11 show interfaces vlan

	show interfaces vlan [<id>] Displays interface configuration.	
Syntax Description	id	Specifies the VLAN ID for which to display data
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.3000	
	3.6.8008	Updated example

	3.9.3100	Updated example to reflect ARP responder and ARP cache-update
Example	switch (config) # show interfaces vlan 100 Vlan 100: Admin state : Enabled Operational state: Down Autostate : Enabled Mac Address : 24:8A:07:83:30:C8 DHCP client : Disabled IPv4 address: 192.168.70.254/24 [primary] 192.168.80.254/24 Broadcast address: 192.168.70.255 [primary] 192.168.80.255 IPv6 address: 4000::1/64 [primary] 5000::1/64 MTU : 1500 bytes Arp timeout : 1500 seconds Arp responder : Disabled Arp cache-update : garp Icmp redirect : Enabled Description : N/A VRF : default Counters : Disabled	
Related Commands		
Notes		

15.1.5.13.12 show ip interface

	show ip interface [vrf <vrf-name>] Displays IP interfaces information.	
Syntax Description	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.6.8008	Updated example
	3.9.3100	Updated example to reflect Arp timeout, Arp responder, Disabled Arp cache-update

Example	switch (config) # show ip interface ethernet 1/1 Eth1/1: Admin state : Enabled Operational state : Up Last change in operational status: 0:11:14 ago (5 oper change) Boot delay time : 0 sec Description : N/A Mac address : 24:8A:07:83:30:C8 MTU : 1500 bytes (Maximum packet size 1522 bytes) Fec : auto Flow-control : receive off send off Supported speeds : 1G 10G 25G Advertised speeds : 1G 10G 25G Actual speed : 25G (auto) Auto-negotiation : Enabled Width reduction mode : Unknown DHCP client : Disabled Autoconfig : Disabled IPv4 address: 192.168.50.254/24 [primary] 192.168.60.254/24 Broadcast address: 192.168.50.255 [primary] 192.168.60.255 IPv6 address: 2000::1/64 [primary] 3000::1/64 fe80::268a:7ff:fe83:30c8/64 Arp timeout : 1500 seconds Arp responder : Disabled Arp cache-update: garp VRF : default Forwarding mode : inherited cut-through …
Related Commands	
Notes	

15.1.5.13.13 show ip interface brief

	show ip interface <vrf vrf_name> brief Displays IP interfaces' brief information.	
Syntax Description	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.6.8008	Updated example
	3.9.0300	Updated example
Example		

```
switch (config) # show  ip interface vrf default brief

-------------------------------------------------------------------------------------------------
Interface        Address/Mask       Primary      Admin-state    Oper-state    MTU          VRF
-------------------------------------------------------------------------------------------------
mgmt0            10.209.21.18/22                 Enabled         Up           1500       default
Loopback 1       1.1.1.1/32         primary      Enabled         Up           1500       default
vrf-default      1.1.1.1/32         primary      Enabled         Up           1500       default
```

Related Commands	
Notes	

15.1.5.13.14 show interfaces configured

	show interfaces [<type> <id>] configured Displays interface configuration.	
Syntax Description	<type> <id>	Specifies the interface for which to display data
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.6.8008	Updated example

Example	`switch (config) # show interfaces mgmt0 configured`
	
	`Interface mgmt0 configuration:`
	`  Comment : `
	`  Enabled : yes`
	`  DHCP : yes`
	`  DHCP Hostname : yes`
	`  Zeroconf : no`
	`  IP address :`
	`  Netmask :`
	`  IPv6 enabled : yes`
	`  Autoconf enabled: no`
	`  Autoconf route : yes`
	`  Autoconf privacy: no`
	`  DHCPv6 enabled : yes`
	`  IPv6 addresses : 0`
	`  Speed : auto`
	`  Duplex : auto`
	`  MTU : 1500`	

Related Commands		
Notes		

15.1.5.13.15 show ip

	show ip interface [vrf <vrf-name>] ethernet <slot>/<port> Displays information on the specified Ethernet interface in the routing-context VRF.	
Syntax Description	<slot>/<port>	Port number
	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.6.8008	Updated example
Example		

```
 switch (config) # show ip interface ethernet 1/1

Eth1/1:
  Admin state                       : Enabled
  Operational state                 : Up
  Last change in operational status: 0:11:14 ago (5 oper change)
  Boot delay time                   : 0 sec
  Description                       : N/A
  Mac address                       : 24:8a:07:83:30:c8
  MTU                               : 1500 bytes (Maximum packet size 1522 bytes)
  Fec                               : auto
  Flow-control                      : receive off send off
  Supported speeds                  : 1G 10G 25G
  Advertised speeds                 : 1G 10G 25G
  Actual speed                      : 25G (auto)
  Auto-negotiation                  : Enabled
  Width reduction mode              : Unknown
  DHCP client                       : Disabled
  Autoconfig                        : Disabled

  IPv4 address:
    192.168.50.254/24 [primary]
    192.168.60.254/24

  Broadcast address:
    192.168.50.255 [primary]
    192.168.60.255

  IPv6 address:
    2000::1/64 [primary]
    3000::1/64
    fe80::268a:7ff:fe83:30c8/64

  Arp responder  : Disabled
  Arp timeout    : 1500 seconds
  VRF            : default
  Forwarding mode: inherited cut-through

  Telemetry sampling: Disabled   TCs: N\A
  Telemetry threshold: Disabled  TCs: N\A
  Telemetry threshold level: N\A

  Last clearing of "show interface" counters: Never
  60 seconds ingress rate                   : 56 bits/sec, 7 bytes/sec, 1 packets/sec
  60 seconds egress rate                    : 8 bits/sec, 1 bytes/sec, 0 packets/sec

  Rx:
    698                   packets
    0                     unicast packets
    0                     multicast packets
    698                   broadcast packets
    44672                 bytes
    0                     discard packets
    0                     error packets
    0                     fcs errors
    0                     undersize packets
    0                     oversize packets
    0                     pause packets
    0                     unknown control opcode
    0                     symbol errors
 Tx:
    1923                  packets
    0                     unicast packets
    1859                  multicast packets
    64                    broadcast packets
    142718                bytes
    0                     discard packets
    0                     error packets
    0                     hoq discard packets
```

Related Commands	
Notes	

15.1.5.13.16 show ip interface mgmt0

	show ip interface [vrf <vrf-name>] mgmt0 Displays management interface information.	
Syntax Description	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.6.8008	Updated example

Example	switch (config) # show ip interface mgmt0
	
	Interface mgmt0 status:
	  Comment :
	  Admin up : yes
	  Link up : yes
	  DHCP running : yes
	  IP address : 10.12.67.33
	  Netmask : 255.255.255.128
	  IPv6 enabled : yes
	  Autoconf enabled: no
	  Autoconf route : yes
	  Autoconf privacy: no
	  DHCPv6 running : yes (but no valid lease)
	  IPv6 addresses : 1
	
	  IPv6 address:
	  fe80::268a:7ff:fe53:3d8e/64
	
	  Speed : 1000Mb/s (auto)
	  Duplex : full (auto)
	  Interface type : ethernet
	  Interface source: bridge
	  MTU : 1500
	  HW address : 24:8A:07:53:3D:8E
	
	  Rx:
	  1843422 bytes
	  25627 packets
	  0 mcast packets
	  0 discards
	  0 errors
	  0 overruns
	  0 frame
	
	  Tx:
	  236174 bytes
	  1897 packets
	  0 discards
	  0 errors
	  0 overruns
	  0 carrier
	  0 collisions
	  0 queue len	

Related Commands		
Notes		

15.1.5.13.17 show ip interface port-channel

	show ip interface [vrf <vrf-name>] port-channel <id> Displays information on the specified LAG in the routing-context VRF.	
Syntax Description	id	LAG ID

	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.6.8008	Updated example
	3.7.1000	Updated example
	3.9.3100	Updated example to reflect Arp timeout, Arp responder, Arp cache-update

Example

```
switch (config) # show ip interface port-channel 1

Po1:
Admin state         : Enabled
Operational state   : Down
Description         : N/A
Mac address         : 24:8A:07:83:30:C8
MTU                 : 1500 bytes (Maximum packet size 1522 bytes)
lacp-individual mode: Disabled
Flow-control        : receive off send off
Actual speed        : 25G (auto)
Auto-negotiation    : N/A
Width reduction mode: Not supported
DHCP client         : Disabled
Autoconfig          : Disabled

IPv4 address:
   192.168.100.254/24 [primary]
   192.168.110.254/24


Broadcast address:
   192.168.100.255 [primary]
   192.168.110.255


IPv6 address:
   6000::1/64 [primary]
   7000::1/64


Arp timeout     : 1500 seconds
Arp responder   : Disabled
Arp cache-update: garp
VRF             : default
Forwarding mode: inherited cut-through
…
```

Related Commands	
Notes	

15.1.5.13.18 show ip interface vrf

	show ip interface vrf {<vrf-name> | all | ethernet <slot>/<port> | loopback <id> | port-channel <id> | vlan <vid>} [brief] Displays IP interface information per VRF.	
Syntax Description	vrf	Displays IP interface information per VRF
	all	Displays information on all VRF
	ethernet	Displays Ethernet interface information per VRF

	loopback	Displays loopback interface information per VRF
	port-channel	Displays LAG information per VRF
	vlan	Displays VLAN interface information per VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.6.5000	Updated example
	3.6.6000	Updated example
	3.6.8008	Updated example
	3.7.1000	Updated example
Example	switch (config) # show ip interface vrf default port-channel 1 Po1: Admin state : Enabled Operational state : Down Description : N/A Mac address : 24:8a:07:83:30:c8 MTU : 1500 bytes (Maximum packet size 1522 bytes) lacp-individual mode: Disabled Flow-control : receive off send off Actual speed : 25G (auto) Auto-negotiation : N/A Width reduction mode: Not supported DHCP client : Disabled Autoconfig : Disabled ...	
Related Commands		
Notes	If no routing-context is specified, the “routing-context” VRF is automatically displayed.	

15.1.5.13.19 show ip interface vrf vrf

	show ip interface vrf <vrf-name> vrf Displays VRF loopback information for a specific VRF.	
Syntax Description	vrf-name	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.9.0300	
Example	switch (config) # show ip interface vrf default	
Related Commands	show ip interface vrf	
Notes		

15.1.5.13.20 show ipv6 interface

	show ipv6 interface Displays IPv6 interface information.	
Syntax Description	vrf	VRF name

Default	N/A
Configuration Mode	Any command mode
History	3.6.8008

Example	`switch (config) # show ipv6 interface`
	
	`Eth1/1:`
	`  VRF        : default`
	`  Admin state: enabled`
	`  IPv6       : enabled`
	
	`  IPv6 address:`
	`    2000::1/64 [primary]`
	`    3000::1/64`
	
	`  Local Link Address:`
	`    fe80::268a:7ff:fe83:30c8/64`
	
	`  Joined group address:`
	`    ff02::1:ff00:1`
	
	`  ND retransmit interval (usec): 1000`
	`  ND DAD                        : enabled`
	`  Number of DAD attempts        : 1`
	`  ND reachable time             : 0`
	
	`Po1:`
	`  VRF        : default`
	`  Admin state: enabled`
	`  IPv6       : enabled`
	
	`  IPv6 address:`
	`    6000::1/64 [primary]`
	`    7000::1/64`
	
	`ND retransmit interval (usec): 1000`
	`  ND DAD                        : enabled`
	`  Number of DAD attempts        : 1`
	`  ND reachable time             : 0`
	
	`vlan100:`
	`  VRF        : default`
	`  Admin state: enabled`
	`  IPv6       : enabled`
	` IPv6 address:`
	`    4000::1/64 [primary]`
	`    5000::1/64`
	
	`  ICMPv6 redirect               : enabled`
	`  ND retransmit interval (usec): 1000`
	`  ND DAD                        : enabled`
	`  Number of DAD attempts        : 1`
	`  ND reachable time             : 0`
	
	`loopback1:`
	`  VRF        : default`
	`  Admin state: enabled`
	`  IPv6       : enabled`
	
	`  IPv6 address:`
	`    2001::1/128 [primary]`
	`    2002::1/128`
	
	`  Local Link Address:`
	`    fe80::4c01:40ff:feb3:b753/64`
	
	`  Joined group address:`
	`    ff02::1:ff00:1`

Related Commands	
Notes	

15.1.5.13.21 show ipv6 interface brief

	show ipv6 interface [vrf <vrf-name>] brief Displays IPv6 interface information.	
Syntax Description	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	
Example		

```
 switch (config) # show ipv6 interface brief
-------------------------------------------------------------------------------------------------------------
---------
Interface     Address/Mask                   Primary    Address-state   Admin-state    Oper-state   MTU
VRF
-------------------------------------------------------------------------------------------------------------
---------
mgmt0         fe80::268a:7ff:fe53:3d8e/64               valid           Enabled        Up           1500
default
mgmt1         fe80::268a:7ff:fe53:3d8f/64               valid           Enabled        Up           1500
default
Eth1/1        2000::1/64                     primary    valid           Enabled        Up           1500
default
Eth1/1        3000::1/64                                valid
Eth1/1        fe80::268a:7ff:fe83:30c8/64               valid
Po1           6000::1/64                     primary    valid           Enabled        Down         1500
default
Po1           7000::1/64                                valid
vlan100       4000::1/64                     primary    valid           Enabled        Down         1500
default
vlan100       5000::1/64                                valid
loopback1     2001::1/128                    primary    valid           Enabled        Up           1500
default
loopback1     2002::1/128                               valid
loopback1     fe80::4c01:40ff:feb3:b753/64              valid
```

Related Commands		
Notes		

15.1.5.13.22 show ipv6

	show ipv6 interface [vrf <vrf-name>] ethernet <slot>/<port> Display IPv6 information of the specified Ethernet interface.	
Syntax Description	<slot>/<port>	Port number
	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	

Example	switch (config) # show ipv6 interface ethernet 1/1
	
	Eth1/1:
	  VRF : default
	  Admin state: enabled
	  IPv6 : enabled
	
	  IPv6 address:
	  2000::1/64 [primary]
	  3000::1/64
	
	  Local Link Address:
	  fe80::268a:7ff:fe83:30c8/64
	
	  Joined group address:
	  ff02::1:ff00:1
	
	  ND retransmit interval (usec): 1000
	  ND DAD : enabled
	  Number of DAD attempts : 1
	  ND reachable time : 0

Related Commands	
Notes	

15.1.5.13.23 show ipv6 interface loopback

	show ipv6 interface [vrf <vrf-name>] loopback <id> Display IPv6 information of the specified loopback interface.	
Syntax Description	id	Loopback port ID
	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	

Example	switch (config) # show ipv6 interface loopback 1
	
	loopback1:
	  VRF : default
	  Admin state: enabled
	  IPv6 : enabled
	
	  IPv6 address:
	  2001::1/128 [primary]
	  2002::1/128
	
	  Local Link Address:
	  fe80::4c01:40ff:feb3:b753/64
	
	  Joined group address:
	  ff02::1:ff00:1	

Related Commands		
Notes		

15.1.5.13.24 show ipv6 interface port-channel

	show ipv6 interface [vrf <vrf-name>] port-channel <id> Display IPv6 information of the specified LAG interface.	
Syntax Description	id	LAG ID

	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	

Example	switch (config) # show ipv6 interface port-channel 1
	
	Po1:
	  VRF : default
	  Admin state: enabled
	  IPv6 : enabled
	
	  IPv6 address:
	  6000::1/64 [primary]
	  7000::1/64
	
	  ND retransmit interval (usec): 1000
	  ND DAD : enabled
	  Number of DAD attempts : 1
	  ND reachable time : 0	

Related Commands		
Notes		

15.1.5.13.25 show ipv6 interface vlan

	show ipv6 interface [vrf <vrf-name>] vlan <vid> Display IPv6 information of the specified VLAN interface.	
Syntax Description	vid	VLAN ID
	vrf	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	

Example	switch (config) # show ipv6 interface vlan 100
	
	vlan100:
	  VRF : default
	  Admin state: enabled
	  IPv6 : enabled
	
	  IPv6 address:
	  4000::1/64 [primary]
	  5000::1/64
	
	  ICMPv6 redirect : disabled
	  ND retransmit interval (usec): 1000
	  ND DAD : enabled
	  Number of DAD attempts : 1
	  ND reachable time : 0	

Related Commands		
Notes		

15.1.5.13.26 show ipv6 interface vrf

	show ipv6 interface vrf <vrf-name> Display IPv6 information of the specified VRF.

Syntax Description	name	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	

Example	switch (config) # show ipv6 interface vrf default
	
	Eth1/1:
	  VRF : default
	  Admin state: enabled
	  IPv6 : enabled
	...
	Po1:
	  VRF : default
	  Admin state: enabled
	  IPv6 : enabled
	...
	vlan100:
	  VRF : default
	  Admin state: enabled
	  IPv6 : enabled
	...
	loopback1:
	  VRF : default
	  Admin state: enabled
	  IPv6 : enabled
	...	

Related Commands		
Notes		

15.1.5.13.27 show ipv6 interface vrf brief

	show ipv6 interface vrf <name> brief Display IPv6 information of the specified VRF in brief form.	
Syntax Description	name	VRF name
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008	
Example		

```
switch (config) # show ipv6 interface vrf default brief
-------------------------------------------------------------------------------------------------------------------
-------
Interface     Address/Mask                  Primary    Address-state   Admin-state   Oper-state   MTU
VRF
-------------------------------------------------------------------------------------------------------------------
-------
mgmt0         fe80::268a:7ff:fe53:3d8e/64              valid          Enabled       Up           1500
default
mgmt1         fe80::268a:7ff:fe53:3d8f/64              valid          Enabled       Up           1500
default
Eth1/1        2000::1/64                    primary    valid          Enabled       Up           1500
default
Eth1/1        3000::1/64                               valid
Eth1/1        fe80::268a:7ff:fe83:30c8/64              valid
Po1           6000::1/64                    primary    valid          Enabled       Down         1500
default
Po1           7000::1/64                               valid
vlan100       4000::1/64                    primary    valid          Enabled       Down         1500
default
vlan100       5000::1/64                               valid
loopback1     2001::1/128                   primary    valid          Enabled       Up           1500
default
loopback1     2002::1/128                              valid
loopback1     fe80::4c01:40ff:feb3:b753/64             valid
```

Related Commands	
Notes	

15.1.5.14 Loopback Interface

15.1.5.14.1 interface loopback

	interface loopback <id> no interface loopback <id> Creates a loopback interface and enters the interface configuration mode. The no form of the command deletes the interface.	
Syntax Description	id	Range: 0-31
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Example	`switch (config) # interface loopback 10` `switch (config interface loopback 10) #`	
Related Commands		
Notes	• Up to 32 loopback interfaces can be configured • Within the loopback configuration mode, you can configure description and ip-address • MTU cannot be configured on the loopback interface	

15.1.5.14.2 interface vrf ip address alias

	interface vrf <vrf-name> ip address alias <loopback<N> | loopback N> no interface vrf <vrf-name> ip address alias Copies addresses from given loopback interface. The no form of the command disables the copied addresses from given loopback interface.

Syntax Description	vrf-name	VRF name
	loopback<N> | loopback N	Loopback interface with specified loopback number
Default	Disabled	
Configuration Mode	config	
History	3.9.0300	
Example	`switch (config)# interface vrf vrf-default ip address alias loopback1`	
Related Commands	show ip interface [vrf]	
Notes		

15.1.5.14.3 ip address

	ip address <ip-address> <mask> no ip address [<ip-address> [<mask>]] Enters user-defined IPv4 address for the interface. The no form of the command removes the specified IPv4 address. If no address is specified, then all IPv4 addresses of this interface are removed.	
Syntax Description	ip-address	IPv4 address
	mask	There are two possible ways to the mask: • /length - only /32 is possible • Network address (i.e. 255.255.255.0) The mask length may be configured without a space (i.e. <ipv4-address>/<length>).
Default	0.0.0.0/0	
Configuration Mode	config	
History	3.3.5006	
Example	`switch (config interface loopback 10) # ip address 10.10.10.10 /32`	
Related Commands	interface loopback	
Notes	An interface may have up to 16 IPv4 address assignments.	

15.1.5.14.4 description

	description <string> no description Enters a description for the interface. The no form of the command sets the description to default.	
Syntax Description	string	User defined string
	mask	There are two possible ways to the mask: • /length - only /32 is possible • Network address (i.e. 255.255.255.0) The mask length may be configured without a space (i.e. <ipv4-address>/<length>).
Default	“”	
Configuration Mode	config interface loopback	

History	3.3.5006
Example	`switch (config interface loopback 10) # description my-ip-interface`
Related Commands	interface loopback
Notes	

15.1.5.14.5 show interfaces loopback

	show interface loopback <id> Displays the attribute of the interface loopback.	
Syntax Description	id	Range: 1-32
Default	N/A	
Configuration Mode	config interface loopback	
History	3.2.3000	
	3.6.8008	Updated example

Example	`switch (config) # show interfaces loopback 1`
	
	`Loopback 1:`
	` IPv4 address:`
	`  192.168.1.1/32 [primary]`
	`  192.168.2.1/32`
	
	` Broadcast address:`
	`  192.168.1.1 [primary]`
	`  192.168.2.1`
	
	` IPv6 address:`
	`  2001::1/128 [primary]`
	`  2002::1/128`
	`  fe80::4c01:40ff:feb3:b753/64`
	
	` MTU : 1500 bytes`
	` Description: N/A`
	` VRF : default`	

Related Commands	interface loopback	
Notes		

15.1.5.15 Routing and ECMP

15.1.5.15.1 ip route

	ip route [vrf <vrf-name>] <ip-prefix> <netmask> {<next -hop-ip-address> | null0} [<distance>] no ip route [vrf <vrf-name>] <ip-prefix> <netmask> [<next -hop-ip-address>] Configures a static route inside VRF. The no form of the command removes the static route configured.	
Syntax Description	vrf-name	VRF session name
	ip-prefix	IP address
	netmask	There are two possible ways to input the mask: • /<length> (e.g. /24) • Network address (e.g. 255.255.255.0)

	next-hop-ip-address	IP address of the next hop
	null0	Sets a static drop-route
	distance	Administrative distance assigned to route. Options include: • No parameter - route is assigned a default administrative distance of 1 • 1-255 - the administrative distance assigned to route
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.4.2008	Added VRF parameter
	3.9.1600	Removed ethernet, port-channel, and vlan parameters
Example	`switch (config) # ip route vrf my-vrf 80.80.80.0 /24 20.20.20.2`	
Related Commands		
Notes	If no routing-context is specified, the “routing-context” VRF is automatically configured.	

15.1.5.15.2 ip load-sharing

	ip load-sharing <type> [ecmp-group-size <size> [max-ecmp-groups <max>]] no ip load-sharing This command sets the ECMP load sharing mode. The no form of the command sets the load-sharing to default.	
Syntax Description	type	• source-ip-port - source ip and TCP/UDP port • destination-ip-port - destination ip and TCP/UDP port • source-destination-ip-port - source & destination ip and TCP/UDP port • flow-label - flow label • udk - user-defined keys • all - all options • consistent - consistent hashing mode
	ecmp-group-size	Configures ECMP consistent hashing group size
	max-ecmp-groups	Configures max groups of ECMP consistent hashing
Default	all	
Configuration Mode	config	
History	3.2.0230	
	3.5.1000	Added flow-label parameter
	3.7.1100	Updated syntax
Example	`switch (config) # ip load-sharing all switch (config) # ip load-sharing consistent [ecmp-group-size<size>]`	
Related Commands	ip route	

Notes	If no routing-context is specified, the “routing-context” VRF is automatically configured.

15.1.5.15.3 show ip route

	show ip route [vrf <vrf-name] [[<ip-address> | <ip-address>/<length>] [longer-prefixes]] [connected | bgp | static] Displays routing table.	
Syntax Description	ip-address	Performs longest prefix match (LPM) and displays best route
	<ip-address>/<length>	Displays next hop for the specified network. If the network does not exist in routing table, it is not shown. Note: It is the user’s responsibility to calculate the mask and enter it correctly. For example: • Valid - show ip route 10.10.10.0/24 • Invalid - show ip route 10.10.10.10/24
	longer-prefixes	Displays the routes to the specified destination and any routes to a more specific destination. (Only available if both IP and mask are specified.)
	connected	Displays entries for routes to networks directly connected to the switch
	bgp	Display BGP routes
	static	Displays entries added through CLI commands
Default	N/A	
History	3.6.5000	Updated example
	3.6.6000	Updated example
	3.6.8008	Updated example
	3.7.1100	Updated example

Example

```
switch (config) # show ip route

Flags:
  F: Failed to install in H/W
  B: BFD protected (static route)
  i: BFD session initializing (static route)
  x: protecting BFD session failed (static route)
  c: consistent hashing
  p: partial programming in H/W

VRF Name default:
  ------------------------------------------------------------------------------
  Destination     Mask            Flag   Gateway         Interface      Source   AD/M
  ------------------------------------------------------------------------------
  default         0.0.0.0                 10.12.67.126    mgmt0          DHCP     1/1
  10.12.67.0      255.255.255.128         0.0.0.0         mgmt0          direct   0/0
  192.168.2.0     255.255.255.0     c     0.0.0.0         vlan1          direct   0/0
```

Related Commands	ip route

Notes	• If no default route exists, then the message “Route not found” is printed
	• Route next hop is BFD controlled, status is viewable when <all> is inserted in the command, and it will be shown as follows:
	  • If route is removed from routing decision it will be marked as “Active”
	  • Protected next hops are marked with “B”
	  • BFD protected failed/non active neighbors are marked with “BF”
	• If no routing-context is specified, the “routing-context” VRF is automatically displayed

15.1.5.15.4 show ip route vrf

	show ip route vrf {<vrf-name> | all} Displays routing table of VRF instance.	
Syntax Description	all	Displays routing tables for all VRF instances
	vrf-name	Name of VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.6.4070	Added support for BFD and updated notes
	3.6.5000	Updated example
	3.6.8008	Updated example

Example

```
switch (config) # show ip route vrf default

Flags:
  F: Failed to install in H/W
  B: BFD protected (static route)
  i: BFD session initializing (static route)
  x: protecting BFD session failed (static route)

VRF Name default:
  ------------------------------------------------------------------------------
  Destination     Mask            Flag   Gateway         Interface      Source   AD/M
  ------------------------------------------------------------------------------
  default         0.0.0.0                 10.12.67.126    mgmt0          DHCP     1/1
  10.12.67.0      255.255.255.128         0.0.0.0         mgmt0          direct   0/0
switch (config) # show ip route vrf my-vrf static

Flags:
  F: Failed to install in H/W
  B: BFD protected (static route)
  i: BFD session initializing (static route)
  x: protecting BFD session failed (static route)

VRF Name my-vrf:
  ------------------------------------------------------------------------------
  Destination     Mask            Flag   Gateway         Interface      Source   AD/M
  ------------------------------------------------------------------------------
  80.80.80.0      255.255.255.0           20.20.20.2      vlan20         static   1/1
```

Related Commands	ip route

Notes	• If no default route exists, then the message “Route not found” is printed
	• Route next hop is BFD controlled, status is viewable when <all> is inserted in the command, and it will be shown as follows:
	  • If route is removed from routing decision it will be marked as “Active”
	  • Protected next hops are marked with “B”
	  • BFD protected failed/non active neighbors are marked with “BF”
	• If no routing-context is specified, the “routing-context” VRF is automatically displayed
	• When using a network prefix, the user must calculate the host mask and enter correctly. For example, “show ip route 10.10.10.0/24” is valid, but “ip route 10.10.10.10/24” is invalid.

15.1.5.15.5 show ip route -a

	show ip route [vrf {<vrf-name> | all}] -a Displays routing table of VRF instance.	
Syntax Description	vrf-name	Name of VRF
	all	Displays routing tables for all VRF instances
	-a	Displays static routes currently inactive due to the interface being down
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.0000	
Example		

```
switch (config) # show ip route vrf my-vrf -a
VRF Name:         my-vrf
-----------------------------
Destination       Mask              Gateway           Interface         Source            Distance/Metric
90.90.90.0        255.255.255.0     1.1.1.2           NA                static            1/0
```

Related Commands	ip route

Notes	• If no default route exists, then the message “Route not found” is printed
	• Route next hop is BFD controlled, status is viewable when <all> is inserted in the command, and it will be shown as follows:
	  • If route is removed from routing decision it will be marked as “Active”
	  • Protected next hops are marked with “B”
	  • BFD protected failed/non active neighbors are marked with “BF”
	• If no routing-context is specified, the “routing-context” VRF is automatically displayed

15.1.5.15.6 show ip route failed

	show ip route [vrf {<vrf-name> | all}] failed Displays failed routes of VRF instance.	
Syntax Description	vrf-name	Name of VRF
	all	Displays routing tables for all VRF instances
Default	N/A	
Configuration Mode	Any command mode	

History	3.6.6000	
	3.6.8008	Updated example

Example

```
switch (config) # show ip route failed

Flags:
  F: Failed to install in H/W
  B: BFD protected (static route)
  i: BFD session initializing (static route)
  x: protecting BFD session failed (static route)

Warning: Number of HW failed routes is 2
These routes are marked with 'f' flag

VRF Name default:
  ------------------------------------------------------------------------------
  Destination     Mask            Flag   Gateway         Interface      Source   AD/M
  ------------------------------------------------------------------------------
  20.20.20.0      255.255.255.0   f      0.0.0.0         vlan20         direct   0/0
  80.80.80.0      255.255.255.0   f      20.20.20.2      vlan20         static   1/1
```

Related Commands	ip route

Notes	• If no default route exists, then the message “Route not found” is printed
	• Route next hop is BFD controlled, status is viewable when <all> is inserted in the command, and it will be shown as follows:
	  ◦ If route is removed from routing decision it will be marked as “Active”
	  ◦ Protected next hops are marked with “B”
	  ◦ BFD protected failed/non active neighbors are marked with “BF”
	• If no routing-context is specified, the “routing-context” VRF is automatically displayed

15.1.5.15.7 show ip route static

	show ip route [vrf {<vrf-name> | all}] static Displays static routes of VRF instance.	
Syntax Description	vrf-name	Name of VRF
	all	Displays routing tables for all VRF instances
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.6.5000	Updated example
	3.6.8008	Updated example

Example

```
switch (config) # show ip route static

Flags:
  F: Failed to install in H/W
  B: BFD protected (static route)
  i: BFD session initializing (static route)
  x: protecting BFD session failed (static route)

VRF Name default:
  ------------------------------------------------------------------------------
  Destination     Mask            Flag   Gateway         Interface      Source   AD/M
  ------------------------------------------------------------------------------
  80.80.80.0      255.255.255.0          20.20.20.2      vlan20         static   1/1
```

Related Commands	ip route

Notes	• If no default route exists, then the message “Route not found” is printed
	• Route next hop is BFD controlled, status is viewable when <all> is inserted in the command, and it will be shown as follows:
	  • If route is removed from routing decision it will be marked as “Active”
	  • Protected next hops are marked with “B”
	  • BFD protected failed/non active neighbors are marked with “BF”
	• If no routing-context is specified, the “routing-context” VRF is automatically displayed

15.1.5.15.8 show ip route static multicast-override

	show ip route [vrf {all | <vrf-name>}] static multicast-override Displays Reverse Path Forwarding (RPF) information for a specific IPv4 multicast source configured via the command “ip mroute”.	
Syntax Description	vrf-name	Name of VRF
	all	Displays information for all VRFs
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6000	
	3.6.8008	Updated example
Example		

```
switch (config) # show ip route vrf default static multicast-override

VRF "default":
  ----------------------------------------------------------------------
  Destination       Mask              Gateway           Route preference
  ----------------------------------------------------------------------
  50.50.50.0        255.255.255.0     20.20.20.45       1
  100.100.8.0       255.255.255.0     20.20.20.9        1
  100.100.100.0     255.255.255.0     20.20.20.22       7
  100.100.100.100   255.255.255.255   20.20.20.9        1
```

Related Commands	
Notes	

15.1.5.15.9 show ip route summary

	show ip route [vrf {<vrf-name> | all}] summary Displays route summary of VRF instance.	
Syntax Description	vrf-name	Name of VRF
	all	Displays routing tables for all VRF instances
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.6.5000	Updated example

Example	switch (config) # show ip route vrf my-vrf summary VRF Name: default ------------------------ Route Source Routes ------------------------ direct 3 static 0 ospf 0 bgp 0 DHCP 1 Total 4
Related Commands	ip route

Notes	• If no default route exists, then the message “Route not found” is printed
	• Route next hop is BFD controlled, status is viewable when <all> is inserted in the command, and it will be shown as follows:
	  • If route is removed from routing decision it will be marked as “Active”
	  • Protected next hops are marked with “B”
	  • BFD protected failed/non active neighbors are marked with “BF”
	• If no routing-context is specified, the “routing-context” VRF is automatically displayed

15.1.5.15.10 show ip route interface

	show ip route [vrf {<vrf-name> | all}] interface {ethernet <slot>/<port> | port-channel <lag> | vlan <vlan>} Displays routing table for specific interfaces.	
Syntax Description	ethernet	Displays routing table for Ethernet interfaces
	port-channel	Displays routing table for LAG interfaces
	vlan	Displays routing table for VLAN interfaces
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	Added VRF parameter
	3.6.5000	Updated example

Example

```
switch (config) # show ip route interface vlan 10
VRF Name:         default
Total number of entries: 1

--------------------------------------------------------------------------------
  Address              Type            Hardware Address          Interface
--------------------------------------------------------------------------------
  15.0.0.2             Static ETH      DE:DE:BE:EF:DE:AD         vlan 10
```

Related Commands	ip route
Notes	

15.1.5.15.11 show ip load-sharing

	show ip load-sharing Displays ECMP hash attribute.
Syntax Description	N/A

Default	N/A	
Configuration Mode	Any command mode	
History	3.4.2008	
	3.7.1100	Updated example

Example	`(config) # show ip load-sharing`
	`Load sharing: all`
	`Type: static`
	
	`(config) # show ip load-sharing`
	`Load sharing: destination-ip-port`
	`Type: consistent`
	`Operational state: stable`
	`Container size: 512`
	`Max number of containers: 40`
	`Used containers: 5`	

Related Commands	ip load-sharing	
Notes	The command's output is different for static & consistent hashing	

15.1.5.16 Network to Media Resolution (ARP)

15.1.5.16.1 ip arp

	ip arp [vrf <vrf-name>] <ip-address> <mac-address> no ip arp <ip-address> Configures IP ARP properties of VRF. The no form of the command deletes the static ARP configuration.	
Syntax Description	vrf-name	VRF session name
	IP address	IPv4 address
	mac-address	MAC address (format XX:XX:XX:XX:XX:XX)
Default	N/A	
Configuration Mode	config	
History	3.4.2008	
Example	`switch (config) # ip arp vrf my-vrf 20.20.20.2 aa:bb:cc:dd:ee:ff`	
Related Commands		
Notes	If no routing-context is specified, the "routing-context" VRF is automatically configured.	

15.1.5.16.2 ip arp responder

	ip arp responder Initiates ARP responder functionality.
Syntax Description	N/A
Default	N/A
Configuration Mode	config interface ethernet config interface port-channel config interface vlan

History	3.6.8008
Example	`switch (config interface vlan 10) # ip arp responder`
Related Commands	ip arp show ip arp
Note	

15.1.5.16.3 ip arp timeout

	ip arp timeout <timeout-value> no ip arp timeout Sets the dynamic ARP cache timeout. The no form of the command sets the timeout to default.	
Syntax Description	timeout-value	Time that an entry remains in the ARP cache Range: 240-28800 seconds
Default	1500 seconds	
Configuration Mode	config interface ethernet config interface port-channel config interface vlan	
History	3.2.0230	
	3.5.1000	Updated Note section
Example	`switch (config interface vlan 10) # ip arp timeout 2000`	
Related Commands	ip arp show ip arp	
Note	• This configuration may take up to 5 minutes to take effect • The time interval after which each ARP entry becomes stale may actually vary from 50-150% of the configured value	

15.1.5.16.4 clear ip arp

	clear ip arp [vrf <vrf-name>] [interface <type> | <IP-address>] Clears the dynamic ARP cache for the specific VRF session.	
Syntax Description	vrf-name	VRF session name
	interface	Clears dynamic ARP entries for a interface
	ip-address	Clears dynamic ARP entries for a specific IP address
Default	N/A	
Configuration Mode	config	
History	3.2.0230	
	3.4.2008	Added VRF parameter
Example	`switch (config) # clear ip arp vrf my-vrf`	
Related Commands	ip arp show ip arp	

Notes	If no routing-context is specified, the “routing-context” VRF is automatically configured.

15.1.5.16.5 show ip arp

	show ip arp [vrf [<vrf-name> | all]] [interface <type> | count | timeout] Displays all ARP information for VRF instance.	
Syntax Description	all	Displays all ARP information for all VRF
	interface	Displays all ARP information for specific interface
	count	Displays number of ARPs for specific VRF
	timeout	Displays value of ARP timeout
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.3000	
	3.4.2008	Added VRF parameter
	3.6.5000	Updated example output
	3.8.2000	Added example of "show ip arp timeout"
	3.9.0500	Updated output example: "Flags" column was added

Example

```
switch (config) # show ip arp

Flags:
G: EVPN Default GW

VRF Name default:
Total number of entries: 4

-------------------------------------------------------------------------------------
Address             Type            Flags      Hardware Address          Interface
-------------------------------------------------------------------------------------
10.209.0.1          Dynamic ETH                00:00:5e:00:01:01         mgmt0
10.209.1.53         Dynamic ETH                24:8a:07:b0:2d:10         mgmt0
6.6.6.6             Dynamic EVPN    G          24:8a:07:ca:cd:48         vlan 6
192.168.10.1        Dynamic ETH                24:8a:07:ca:cd:48         eth 1/10
```

Example (show ip arp timeout)

```
switch (config)# show ip arp timeout
---------------------------------------
VRF Timeout(in seconds)
---------------------------------------
vrf-default 1500
```

Related Commands	ip arp
Notes	If no routing-context is specified, the “routing-context” VRF is automatically displayed.

15.1.5.17 IP Diagnostic Tools

15.1.5.17.1 ping

	ping [vrf <vrf-name>] [-LRUbdfnqrvVaA] [-c count] [-i interval] [-w deadline] [-p pattern] [-s packetsize] [-t ttl] [-I interface or address] [-M mtu discovery hint] [-S sndbuf] [-T timestamp option] [-Q tos] [hop1 ...] destination Sends ICMP echo requests to a specified host.	
Syntax Description	vrf	Specifies VRF instance name
	Linux Ping options	
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.4.2008	Added VRF parameter
Example		

```
switch (config) # ping 172.30.2.2
PING 172.30.2.2 (172.30.2.2) 56(84) bytes of data.
64 bytes from 172.30.2.2: icmp_seq=1 ttl=64 time=0.703 ms
64 bytes from 172.30.2.2: icmp_seq=2 ttl=64 time=0.187 ms
64 bytes from 172.30.2.2: icmp_seq=3 ttl=64 time=0.166 ms
64 bytes from 172.30.2.2: icmp_seq=4 ttl=64 time=0.161 ms
64 bytes from 172.30.2.2: icmp_seq=5 ttl=64 time=0.153 ms
64 bytes from 172.30.2.2: icmp_seq=6 ttl=64 time=0.144 ms
^C
--- 172.30.2.2 ping statistics ---
6 packets transmitted, 6 received, 0% packet loss, time 5004ms
rtt min/avg/max/mdev = 0.144/0.252/0.703/0.202 ms
```

Related Commands	traceroute
Notes	When using -I option use the interface name + interface number, for example “ping -I vlan10”

15.1.5.17.2 traceroute

	traceroute [vrf <vrf-name>] [-46dFITUnrAV] [-f first_ttl] [-g gate,...] [-i device] [-m max_ttl] [-N squeries] [-p port] [-t tos] [-l flow_label] [-w waittime] [-q nqueries] [-s src_addr] [-z sendwait] host [packetlen] Traces the route packets take to a destination.	
Syntax Description	vrf	Specifies VRF instance name
	-4	Uses IPv4
	-6	Uses IPv6
	-d	Enables socket level debugging
	-F	Sets DF (“do not fragment” bit) on
	-I	Uses ICMP ECHO for tracerouting
	-T	Uses TCP SYN for tracerouting
	-U	Uses UDP datagram (default) for tracerouting
	-n	Does not resolve IP addresses to their domain names

	-r	Bypasses the normal routing and send directly to a host on an attached network
	-A	Performs AS path lookups in routing registries and print results directly after the corresponding addresses
	-V	Prints version info and exit
	-f	Starts from the first_ttl hop (instead from 1)
	-g	Routes packets throw the specified gateway (maximum 8 for IPv4 and 127 for IPv6)
	-i	Specifies a network interface to operate with
	-m	Sets the max number of hops (max TTL to be reached) Default: 30
	-N	Sets the number of probes to be tried simultaneously Default: 16
	-p	Uses destination port. It is an initial value for the UDP destination port (incremented by each probe, default is 33434), for the ICMP seq number (incremented as well, default from 1), and the constant destination port for TCP tries (default is 80).
	-t	Sets the TOS (IPv4 type of service) or TC (IPv6 traffic class) value for outgoing packets
	-l	Uses specified flow_label for IPv6 packets
	-w	Sets the number of seconds to wait for response to a probe (default is 5.0). Non-integer (float point) values allowed too.
	-q	Sets the number of probes per each hop Default: 3
	-s	Uses source src_addr for outgoing packets
	-z	Sets minimal time interval between probes (default is 0). If the value is more than 10, then it specifies a number in milliseconds, else it is a number of seconds (float point values allowed too).
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.4.2008	Added VRF parameter

Example

```
switch (config) # traceroute 192.168.10.70
traceroute to 192.168.10.70 (192.168.10.70), 30 hops max, 40 byte packets
1 172.30.0.1 (172.30.0.1) 3.632 ms 2.849 ms 3.544 ms
2 10.222.128.46 (10.222.128.46) 3.176 ms 3.289 ms 3.656 ms
3 10.158.128.30 (10.158.128.30) 15.331 ms 15.819 ms 16.388 ms
4 10.158.128.65 (10.158.128.65) 20.468 ms 7.893 ms 12.27 ms
5 10.7.34.115 (10.7.34.115) 16.405 ms 11.985 ms 12.264 ms
6 192.168.10.70 (192.168.10.70) 16.377 ms 16.091 ms 20.475 ms
```

Related Commands	
Notes	• The following flags are not supported: -6, -l, -A • When using -i option use the interface name + interface number, for example “traceroute -i vlan10”

15.1.5.17.3 tcpdump

	tcpdump [vrf <vrf-name>] [-aAdeflLnNOpqRStuUvxX] [-c count] [-C file_size] [-E algo:secret] [-F file] [-i interface] [-M secret] [-r file] [-s snaplen] [-T type] [-w file] [-W filecount] [-y datalinktype] [-Z user] [expression] Invokes standard binary, passing command line parameters straight through. Runs in foreground, printing packets as they arrive, until the user hits Ctrl+C.	
Syntax Description	vrf	Specifies VRF instance name
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.4.2008	Added VRF parameter
Example		

```
switch (config) # tcpdump
......
09:37:38.678812 IP 192.168.10.7.ssh > 192.168.10.1.54155: P 1494624:1494800(176) ack 625 win 90
<nop,nop,timestamp 5842763 858672398>
09:37:38.678860 IP 192.168.10.7.ssh > 192.168.10.1.54155: P 1494800:1495104(304) ack 625 win 90
<nop,nop,timestamp 5842763 858672398>
...
9141 packets captured
9142 packets received by filter
0 packets dropped by kernel
```

Related Commands	
Notes	• When using -i option use the interface name + interface number, for example “tcpdump -i vlan10” • For all flag options of this command refer to the linux ‘man page’ of tcp dump

15.1.5.18 QoS

15.1.5.18.1 qos map dscp-to-pcp preserve-pcp

	qos map dscp-to-pcp preserve-pcp no qos map dscp-to-pcp preserve-pcp Configures the router to copy PCP bits when transferring data from one subnet to another. The no form of the command disables this ability.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.3.4000
Example	`switch (config) # qos map dscp-to-pcp preserve-pcp`
Related Commands	
Notes	

15.1.5.19 PBR

15.1.5.19.1 nexthop-group direct

	<ip|ipv6> pbr nexthop-group <group_name> [recursive] no <ip|ipv6> pbr nexthop-group <group_name> Creates direct or recursive nexthop-group and enter to the nexthop-group CLI context. The no form of the command deletes the nexthop-group.	
Syntax Description	group_name	Name of the desired nexthop-group
Default	Disabled	
Configuration Mode	config	
History	3.9.2000	
Example	`switch (config) # ip pbr nexthop-group n_ggg_v4` `switch (config) # ipv6 pbr nexthop-group n_ggg_v6`	
Related Commands	show pbr nexthop-group pbr route-map seq set nexthop-group	
Notes	Maximum number of created nexthop-groups is 1000. Name for the nexthop-group with different IP family also should be different.	

15.1.5.19.2 nexthop-group direct nexthop interface

	<ip|ipv6> pbr nexthop-group <group_name> nexthop interface {ethernet <port/slot>|port-channel <ID> | vlan <ID>} <next-hop IP address> no <ip|ipv6> pbr nexthop-group <group_name> nexthop interface {ethernet <port/slot>|port-channel <ID> | vlan <ID>} <next-hop IP address> Adds nexthop (L3 interface and nexthop IP address) to requested nexthop-group. The no form of the command deletes the desired nexthop from the nexthop-group.	
Syntax Description	group_name	Name of the desired nexthop-group
	port/slot	Physical port
	port-channel <ID>	LAG
	vlan <ID>	VLAN ID
Default	Disabled	
Configuration Mode	config	
History	3.9.2000	

Example	`switch (config) # ip pbr nexthop-group n_ggg_v4 nexthop interface ethernet 1/4 10.10.10.23`
	
	`switch (config) # ipv6 pbr nexthop-group n_ggg_v6 nexthop interface vlan 5 194:23::2`	

Related Commands	show pbr nexthop-group	
Notes	Maximum number of configured direct nexthops in one group is 128. One nexthop can be configured only once in one nexthop-group.	

15.1.5.19.3 nexthop-group recursive nexthop

	<ip|ipv6}>pbr nexthop-group <group_name> recursive nexthop vrf <vrf_name> <next-hop IP address> no <ip|ipv6> pbr nexthop-group <group_name> recursive no nexthop Adds recursive nexthop to requested nexthop-group. The no form of the command deletes the desired nexthop from the nexthop-group.	
Syntax Description	group_name	Name of the desired nexthop-group
	vrf_name	VRF where the desired nexthop is placed
	next-hop IP address	IP/IPv6 address of desired nexthop
Default	N/A	
Configuration Mode	config	
History	3.9.2000	

Example	`switch (config) # ip pbr nexthop-group n_ggg_v4 recursive nexthop vrf default 10.10.10.23`
	
	`switch (config) # ipv6 pbr nexthop-group n_ggg_v6 recursive nexthop vrf default 194:23::2`	

Related Commands	show pbr nexthop-group	
Notes	Maximum number of configured recursive nexthops in one nexthop-group is 1.	

15.1.5.19.4 route-map

	<ip|ipv6> pbr route-map <map_name> no <ip|ipv6> pbr route-map <map_name> Creates route-map and enter to the route-map CLI context. The no form of the command deletes the route-map.	
Syntax Description	map_name	Name of the desired nexthop-group
	vrf_name	VRF where the desired nexthop is placed
	next-hop IP address	IP/IPv6 address of desired nexthop
Default	N/A	
Configuration Mode	config	
History	3.9.2000	

Example	`switch (config) # ip pbr route-map r_ttt_v4`
	
	`switch (config) # ipv6 pbr route-map r_ttt_v6`	

Related Commands	show pbr route-map	
Notes	Maximum number of configured route-maps is 200. Name for the route-map with different IP family also should be different.	

15.1.5.19.5 route-map sequence match rule

	<ip|ipv6> pbr route-map <map_name> seq <number> match {dest-addr <IP address/prefix length>|source-addr <IP address/prefix length>|protocol <tcp|udp>|source-port <port>| dest-port<port>|dscp <value>} no <ip|ipv6> pbr route-map <map_name> seq <number> match Create or modify sequence with new match rule. No form deletes match rule from the sequence.	
Syntax Description	map_name	Name of the desired nexthop-group
	number	ID of sequence inside of the route-map
	IP address/ prefix length	IPv4/IPv6 subnet to be matched on the packet
	tcp|udp	Protocol type to be matched on the packet
	port	Desired TCP or UDP protocol to be matched on the packet
	value	DSCP value to be matched on the packet
Default	N/A	
Configuration Mode	config	
History	3.9.2000	

Example	`switch (config) # ip pbr route-map r_ttt_v4 seq 3 match dest-addr 1.2.3.0/24 source-addr 4.5.6.0/24 dest-port 656 source-port 757 protocol tcp`
	
	`switch (config) # ipv6 pbr route-map r_ttt_v6 seq 3 match dest-addr 23:23::/64 source-addr 90:23::/64 dest-port 656 source-port 757 protocol tcp`	

Related Commands	show pbr route-map	
Notes	Match for source/destination IP address should be specified according to route-map IP family. Maximum number of sequences is 2000 (totally in the system). Maximum number of IPv6 sequences is 2000. Sequence field can be omitted, in this case system with generate new sequence number with index +10 for the last created. Currently DSCP value can be only {0,1,2,3,4}.	

15.1.5.19.6 route-map sequence nexthop-group

	<ip|ipv6> pbr route-map <map_name> seq <number> set nexthop-group <group_name> no <ip|ipv6> pbr route-map <map_name> seq <number> set nexthop-group Specify desired nexthop-group for sending matched traffic. The no form of the command deletes the biding.	
Syntax Description	map_name	Name of the desired nexthop-group
	number	ID of sequence inside of the route-map
	group_name	Name of the desired nexthop-group
Default	N/A	
Configuration Mode	config	
History	3.9.2000	

Example	switch (config) # ip pbr route-map r_ttt_v4 seq 3 set nexthop-group n_ggg_v4 switch (config) # ipv6 pbr route-map r_ttt_v6 seq 3 set nexthop-group n_ggg_v6
Related Commands	show pbr route-map
Notes	One nexthop-group can be specified in more than one sequence.

15.1.5.19.7 route-map sequence counter

	<ip|ipv6> pbr route-map <map_name> seq <number> counter no <ip|ipv6> pbr route-map <map_name> seq <number> counter Request counter to be allocated and count matched packets by match rule. The no form of the command de-allocate the counter.	
Syntax Description	map_name	Name of the desired nexthop-group
	number	ID of sequence inside of the route-map
Default	N/A	
Configuration Mode	config	
History	3.9.2000	
Example	switch (config) # ip pbr route-map r_ttt_v4 seq 3 counter switch (config) # ipv6 pbr route-map r_ttt_v6 seq 3 counter	
Related Commands	show pbr route-map	
Notes		

15.1.5.19.8 bind/unbind route-map on interface

	interface {ethernet <slot/port>|port-channel <ID> | vlan <ID>} {ip|ipv6} pbr route-map <map_name> no interface {ethernet <slot/port>|port-channel <ID> | vlan <ID>} {ip|ipv6} pbr route-map <map_name> Bind requested route-map on the ingress router port. Te no form of the command unbinds requested route-map from the interface.	
Syntax Description	map_name	Name of the desired nexthop-group
	ethernet <port/slot>	Physical port
	port-channel <ID>	LAG
	vlan <ID>	VLAN ID
Default	N/A	
Configuration Mode	config	
History	3.9.2000	

Example	switch (config) # interface ethernet 1/2 ip pbr route-map r_ttt_v4 switch (config) # interface vlan 3 ipv6 pbr route-map r_ttt_v6
Related Commands	show ip interface
Notes	In one time one IPv4 and one IPv6 route-map can be bound on interface

15.1.5.19.9 show nexthop-groups

	show {ip|ipv6} pbr nexthop-group brief | <name> Shows brief information about the all configured nexthop-groups. In case of specifying nexthop-group name show details.	
Syntax Description	name	Name of the desired nexthop-group
Default	N/A	
Configuration Mode	config	
History	3.9.2000	
Example	switch (config) # show ip pbr nexthop-group brief Flags: A: active I: inactive F: failed to install in H/W --------- Name Type Flags Notes --------- n_ggg_v4 direct I Group doesn't have active/ resolved next- hops switch (config) # show ip pbr nexthop-group bbb Flags: A: active I: inactive F: failed to install in H/W bbb: Type : direct Egress interface: vlan 4 (10.10.10.23) Flags : A Notes: N/A	
Related Commands		
Notes	In case of any misconfiguration field “Notes” will reflect it.	

15.1.5.19.10 show route-maps

	show {ip|ipv6} pbr route-map brief |<name> Shows brief information about the all configured route-maps. In case of specifying route-map name show details.	
Syntax Description	name	Name of the desired nexthop-group
Default	N/A	
Configuration Mode	config	

History	3.9.2000
Example	switch (config) # show ip pbr route-map brief Flags: A: active I: inactive F: failed to install in H/W ---------- Name Total sequences Active/Inactive Bound to interfaces r_ttt_v4 1 0/1 switch (config) # show ip pbr route-map r_ttt_v4 Flags: A: active I: inactive F: failed to install in H/W tests: seq match counter nexthop-group flags 1 protocol tcp 0 n_ggg_v4 A
Related Commands	
Notes	In case of any misconfiguration field “Notes” will reflect it.

15.1.5.19.11 route-map to interface bind

	show {ip|ipv6} interface {ethernet <slot/port>| port-channel <ID> | vlan <ID>} Show binding the route-map to interface and its route-map state.	
Syntax Description	ethernet <port/slot>	Physical port
	port-channel <ID>	LAG
	vlan <ID>	VLAN ID
Default	N/A	
Configuration Mode	config	
History	3.9.2000	

Example	switch (config) # show ip interface vlan 3
	
	Vlan 3:
	  Admin state : Enabled
	  Operational state : Down
	  Autostate : Enabled
	  Mac Address : 7c:fe:90:f6:aa:08
	  DHCP client : Disabled
	  PBR route-map : r_ttt_v4
	  PBR route-map state: Active
	…
	
	switch (config) # show ipv6 interface vlan 3
	vlan3:
	  VRF : default
	  Admin state : enabled
	  IPv6 : enabled
	  ICMPv6 redirect : disabled
	  ND retransmit interval (usec): 1000
	  ND DAD : enabled
	  Number of DAD attempts : 1
	  ND reachable time : 0
	  PBR route-map : r_ttt_v6
	  PBR route-map state : Active

Related Commands	
Notes	In case of any misconfiguration field “Notes” will reflect it.

15.1.5.19.12 show pbr general information

	show ip pbr [exceptions [nexthop-group | route-map | interface]] Shows number of configured nexthop-groups and route-maps. Also shows misconfiguration for all PBR configuration or per selected category.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.9.2000

Example	switch (config) # show ip pbr
	
	Configured nexthop-groups (active/inactive/not installed): 1 (0/1/0)
	Configured route-maps (bind/unbind) : 1 (1/0)
	
	switch (config) # show ip pbr exceptions
	
	Configured nexthop-groups (active/inactive/not installed): 1 (0/1/0)
	Configured route-maps (bind/unbind) : 1 (1/0)
	
	Exceptions:
	  Nexthop-groups:
	  Nexthop-group n_ggg_v4 doesn't have active/resolved next-hops
	
	  Route-maps:
	  Route-map r_ttt_v4 sequence 1 assigned nexthop-group is not Active
	
	  Interfaces:
	  Interface vlan 3 assigned route-map is not Active

Related Commands	
Notes	Information about the total amount of configured nexthop-groups and route-maps includes both IPv4 and IPv6 families and does not depend on the specified IP family in CLI command.

15.1.6 IPv6

IP version 6 (IPv6) is a routing protocol which succeeds IPv4. With the expansion of the Internet and databases IPv6 addresses consist of 128 bits whose purpose is to allow networks to include a significantly higher number of nodes by increasing the pool of available unique IP addresses. IPv6 packets alleviate overhead and allow for future customizability.

Textual representations of IPv6 addresses consist of 128 bits made up from eight 16-bit hexadecimal numbers separated by colons. IPv6 addresses may be abbreviated as follows:

- You may omit leading zeros in each 16-bit sequence
- You may replace an entire sequence with a double colon if it equals zero

For example, these addresses represent the same IPv6 address:

- af23:0000:0000:0000:1284:037d:35ce:2401
- af23:0:0:0:1284:37d:35ce:2401
- af23::1284:37d:35ce:2401

IPv6 addresses typically denote a 64-bit network prefix and a 64-bit host address.

15.1.6.1 Features that Support IPv6

The following are the features IPv6 is supported on:

- Static Routes
- ECMP
- Neighbor Discovery
- BGP
- BFD for BGP (IPv6), and Static Routes
- DHCPv6 Relay

15.1.6.2 Neighbor Discovery Protocol

Neighbor Discovery (ND) decides relationships between neighbors and replaces ARP, ICMP, and ICMP redirect in IPv4.

Five kinds of ICMPv6 packets are defined by ND:

- Neighbor advertisement
- Router advertisement
- Neighbor solicitation
- Router solicitation
- Redirect

ND checks whether a neighboring node’s address has changed, whether the neighbor is still reachable, and also resolves the address of the neighbor which a packet is being forwarded to. ND is also useful for network nodes for discovering other nodes and performing basic link-layer configuration.

15.1.6.3 Configuring IPv6

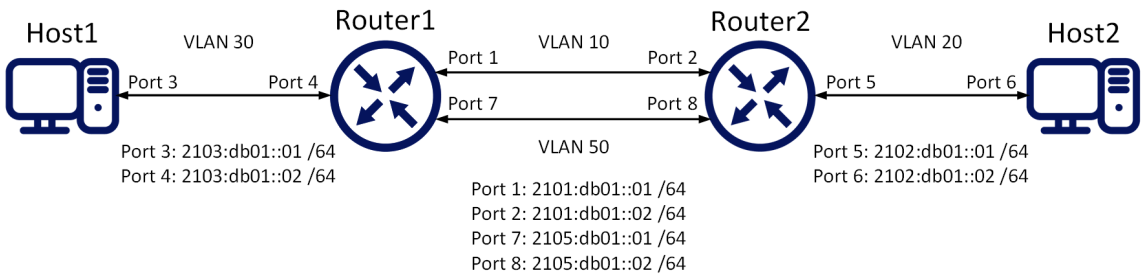

To configure Router1:

1. Enable IP routing. Run:

```
switch (config)# ip routing
```

2. Enable forwarding IPv6 unicast packets. Run:

```
switch (config)# ipv6 routing
```

3. Configure the VLAN interfaces. Run:

```
switch (config)# interface vlan 10
switch (config interface vlan 10) # exit
switch (config)# interface vlan 30
switch (config interface vlan 30) # exit
switch (config)# interface vlan 50
switch (config interface vlan 50) # exit
```

4. Enable IPv6 on the VLAN interfaces. Run:

```
switch (config)# interface vlan 10 ipv6 enable
switch (config)# interface vlan 30 ipv6 enable
switch (config)# interface vlan 50 ipv6 enable
```

5. Configure IPv6 addresses for each one of the VLAN interfaces. Run:

```
switch (config)# interface vlan 10 ipv6 address 2101:db01::1 /64
switch (config)# interface vlan 30 ipv6 address 2103:db01::2 /64
switch (config)# interface vlan 50 ipv6 address 2105:db01::1 /64
```

6. Configure IPv6 unicast on port 2. Run:

```
switch (config)# ipv6 route 2002:db01:: /64 2101:db01::2
```

7. Configure IPv6 unicast on port 8. Run:

```
switch (config)# ipv6 route 2002:db01:: /64 2105:db01::2
```

To configure Router2:

1. Disable prefix mode on the CLI. Run:

```
switch (config)# no cli default prefix-mode enable
```

2. Enable the VLANs on the system. Run:

```
switch (config)# vlan 10
switch (config vlan 10) # exit
switch (config)# vlan 20
switch (config vlan 20) # exit
switch (config)# vlan 50
switch (config vlan 50) # exit
```

3. Configure the switch ports to accept the VLANs of which they are part only. Run:

```
switch (config)# interface ethernet 1/1 switchport access vlan 10  // port2
switch (config)# interface ethernet 1/2 switchport access vlan 50  // port8
switch (config)# interface ethernet 1/36 switchport access vlan 20 // port5
```

4. Disable spanning tree. Run:

```
switch (config)# no spanning-tree
```

5. Enable forwarding IPv6 unicast packets. Run:

```
switch (config)# ipv6 routing
```

6. Configure the VLAN interfaces. Run:

```
switch (config)# interface vlan 10
switch (config interface vlan 10) # exit
switch (config)# interface vlan 20
switch (config interface vlan 20) # exit
switch (config)# interface vlan 50
switch (config interface vlan 50) # exit
```

7. Configure IPv6 addresses for each one of the VLAN interfaces. Run:

```
switch (config)# interface vlan 10 ipv6 address 2101:db01::2 /64
switch (config)# interface vlan 20 ipv6 address 2102:db01::1 /64
switch (config)# interface vlan 50 ipv6 address 2105:db01::2 /64
```

8. Configure IPv6 unicast on port 1. Run:

```
switch (config)# ipv6 route 2103:db01:: /64 2101:db01::1
```

9. Configure IPv6 unicast on port 7. Run:

```
switch (config)# ipv6 route 2103:db01:: /64 2105:db01::1
```

Ping neighbor to verify IPv6 configuration:

```
switch (config)# ping6 2101:db01::2
PING 2101:db01::2(2101:db01::2) 56 data bytes
64 bytes from 2101:db01::2: icmp_seq=1 ttl=64 time=0.371 ms
64 bytes from 2101:db01::2: icmp_seq=2 ttl=64 time=0.620 ms
64 bytes from 2101:db01::2: icmp_seq=3 ttl=64 time=0.192 ms
64 bytes from 2101:db01::2: icmp_seq=4 ttl=64 time=0.277 ms
64 bytes from 2101:db01::2: icmp_seq=5 ttl=64 time=0.231 ms
```

15.1.6.4 IPv6 Commands

15.1.6.4.1 ipv6 enable

	ipv6 enable no ipv6 enable Assigns automatic link-local IPv6 address to the interface. The no form of the command de-assigns that automatic local address and disables IPv6 if no static IPv6 address has been assigned to the interface.	
Syntax Description	N/A	
Default	Unassigned	
Configuration Mode	config interface vlan config interface loopback config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated notes and command description
Example	`switch (config vlan 10) # ipv6 enable`	
Related Commands		
Notes	Assigning an IPv6 address to an interface also enables IPv6 processing on the interface.	

15.1.6.4.2 ipv6 address

	ipv6 address <ipv6-address> /<length> no ipv6 address [<ipv6-address> [/<length>]] Enables IPv6 processing and assigns an IPv6 address to the interface. The no form of the command removes the specified IPv6 address. If no address is specified, then all addresses of the interface are removed.	
Syntax Description	ipv6-address	IPv6 address
	length	Mask length for the associated address space Range: 1-128 The mask length may be configured without a space (i.e. <ipv6-address>/<length>)
Default	N/A	
Configuration Mode	config interface vlan config interface loopback config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated syntax description and example output
Example	`switch (config vlan 10) # ipv6 address 2001::1 /120` `switch (config vlan 10) # ipv6 address 2001::1/120`	
Related Commands		
Notes	An interface may have up to 16 IPv6 address assignments	

15.1.6.4.3 ipv6 nd managed-config-flag

	ipv6 nd managed-config-flag no ipv6 nd managed-config-flag Sets the managed address configuration flag in IPv6 router advertisements. The no form of the command restores the default setting.	
Syntax Description	N/A	
Default	Managed address configuration flag is not set	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated configuration mode
Example	`switch (config vlan 10) # ipv6 nd managed-config-flag`	
Related Commands		
Notes		

15.1.6.4.4 ipv6 nd ns-interval

	ipv6 nd ns-interval <period> no ipv6 nd ns-interval Configures the interval between IPv6 neighbor solicitation (NS) transmissions. The no form of the command restores the default value.	
Syntax Description	period	Time in milliseconds Range: 1000-4294967295
Default	1000	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated configuration mode
Example	`switch (config vlan 10) # ipv6 nd ns-interval 1500`	
Related Commands		
Notes		

15.1.6.4.5 ipv6 nd other-config-flag

	ipv6 nd other-config-flag no ipv6 nd other-config-flag Indicates that other configuration information is available via DHCPv6. The no form of the command removes the other configuration flag.
Syntax Description	N/A
Default	Not set

Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated configuration mode
Example	`switch (config vlan 10) # ipv6 nd other-config-flag`	
Related Commands		
Notes		

15.1.6.4.6 ipv6 nd prefix

	ipv6 nd prefix <ipv6-address> /<length> [no-autoconfig] [no-onlink] [valid-time {<time> | infinite}] [preferred-time {<time> | infinite}] ipv6 nd prefix <prefix> no-advertise no ipv6 nd prefix <prefix> Configures inclusion for router advertisements (RAs) for neighbor. The no form of the command removes the corresponding IPv6 nd prefix.	
Syntax Description	ipv6-address	IPv6 address
	length	Prefix length for the associated address space Range: 1-128
	no-advertise	Prevents advertising of the specified default prefix
	valid-time	Time in seconds Range: 0-4294967295
	preferred-time	Time in seconds Range: 0-4294967295
	no-autoconfig	Indicates that the prefix cannot be used for stateless address configuration
	no-onlink	Indicates that the prefix cannot be used for on-link determination
Default	valid-time: 2592000 seconds preferred-time: 604800 seconds no-autoconfig: Reset, autoconfig enabled no-onlink: Reset, on-link determination is enabled	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated syntax description, configuration mode and default values
Example	`switch (config vlan 10) # ipv6 nd prefix 2001::1 /120`	
Related Commands		
Notes	• Valid time must be larger than preferred time • By default, the router advertises all configured sunbnets on the interface	

15.1.6.4.7 ipv6 nd ra dns-servers lifetime

	ipv6 nd ra dns-servers lifetime {<time> | infinite} no ipv6 nd ra dns-servers lifetime Advertises a lifetime of a Recursive DNS Server (RDNSS). The no form of the command resets the lifetime value to default.	
Syntax Description	time	Possible values: • 0 – RDNSS address can no longer be used • 1-4294967295 in seconds
	infinite	A value of all one bits (0xffffffff) and “infinite” represents infinity
Default	If no lifetime period is configured on the interface, the default value is 1.5 times the Router Advertisement (RA) interval set by the command “ipv6 nd ra interval”	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated command and syntax description, configuration mode and default values
Example	`switch (config vlan 10) # ipv6 nd ra dns-servers lifetime infinite`	
Related Commands		
Notes	• Using the RDNSS and DNSSL options, an IPv6 host can perform IPv6 address network configuration and DNS information simultaneously, without using DHCPv6 for the DNS configuration • A lifetime value set for an individual RDNSS overrides this value • The lifetime value is the maximum amount of time after a route advertisement packet is sent that the RDNSS referenced in the packet may be used for name resolution	

15.1.6.4.8 ipv6 nd ra dns-server

	ipv6 nd ra dns-server <ipv6 address> [lifetime [<time> | infinite]] no ipv6 nd ra dns-server [<ipv6 address>] Configures the IPv6 address of a Recursive DNS Server (RDNSS) to include in the neighbor-discovery router advertisements (RAs). The no form of the command removes the RDNSS from the configuration.	
Syntax Description	ipv6 address	IPv6 address of RDNSS
	lifetime	Maximum lifetime value for the specified RDNSS entry. Possible values: • 0 – RDNSS address can no longer be used • 1-4294967295 in seconds
	infinite	A value of all one bits (0xffffffff) and “infinite” represents infinity
Default	If no lifetime period is configured on the interface, the default value is 1.5 times the Router Advertisement (RA) interval set by the command “ipv6 nd ra interval”	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	

	3.6.4110	Updated command, example and syntax description, configuration mode and default values
Example	`switch (config vlan 10) # ipv6 nd ra dns-server 2001::1 lifetime infinite`	
Related Commands		
Notes	• Including RDNSS information in RAs provides DNS server configuration for connected IPv6 hosts without requiring DHCPv6 • Multiple servers can be configured on the interface by using the command repeatedly • A lifetime value for the RDNSS can optionally be specified with this command, and overrides any default value configured for the interface using the ipv6 nd ra dns-servers lifetime command	

15.1.6.4.9 ipv6 nd ra dns-suffixes lifetime

	ipv6 nd ra dns-suffixes <domain-name> lifetime {<time> | infinite} no ipv6 nd ra dns-suffixes <domain-name> lifetime Advertises a lifetime of a DNS Search List (DNSSL). Using RDNSS and DNSSL options, an IPv6 host can perform IPv6 address network configuration and DNS information simultaneously, without using DHCPv6 for the DNS configuration. The no form of the command resets the lifetime value to its default.	
Syntax Description	time	Possible values: • 0 – RDNSS address can no longer be used • 1-4294967295 in seconds
	infinite	A value of all one bits (0xffffffff) and “infinite” represents infinity
Default	If no lifetime period is configured on the interface, the default value is 1.5 times the Router Advertisement (RA) interval set by the command “ipv6 nd ra interval”	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated command, example and syntax description, configuration mode and default values
Example	`switch (config vlan 10) # ipv6 nd ra dns-suffix domain.com lifetime infinite`	
Related Commands		
Notes	The DNSSL contains the domain names of DNS suffixes for IPv6 hosts to append to short, unqualified domain names for DNS queries	

15.1.6.4.10 ipv6 nd ra dns-suffix

	ipv6 nd ra dns-suffix <domain-name> [lifetime {<time> | infinite}] no ipv6 nd ra dns-suffix [<domain-name>] Creates a DNS search list (DNSSL) to include in the neighbor-discovery Router Advertisements (RAs). The no form of the command removes the DNSSL from the configuration.

Syntax Description	domain-name	Domain suffix for IPv6 hosts to append to short unqualified domain names for DNS queries The suffix must contain only alphanumeric characters, “.” (periods), “-” (hyphens), and must begin and end with an alphanumeric character
	lifetime	Maximum lifetime value for the specified DNSSL entry
	time	Possible values: • 0 – DNSSL must not be used for name resolution • 1-4294967295 in seconds
	infinite	A value of all one bits (0xffffffff) and “infinite” represents infinity
Default	If no lifetime period is configured on the interface, the default value is 1.5 times the Router Advertisement (RA) interval set by the command “ipv6 nd ra interval”	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated command, example and syntax description, configuration mode and default values
Role	admin	
Example	`switch (config vlan 10) # ipv6 nd ra dns-suffix domain.com lifetime infinite`	
Related Commands		
Notes	• The DNSSL contains the domain names of DNS suffixes for IPv6 hosts to append to short, unqualified domain names for DNS queries • Multiple DNS domain names can be added to the DNSSL by reusing the command • A lifetime value for the DNSSL can optionally be specified with this command which overrides any default value configured for the interface using the command “ipv6 nd ra dns-suffixes lifetime”	

15.1.6.4.11 ipv6 nd ra hop-limit

	ipv6 nd ra hop-limit <limit> no ipv6 nd ra hop-limit Sets a suggested hop-limit value to be included in route advertisement (RA) packets. The no form of the command resets the parameter to its default value.	
Syntax Description	limit	The hop-limit value to be included by attached hosts in outgoing packets. • 0 – unspecified (by this router) • 1-255 – number of hops
Default	Limit value is 64	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated configuration modes
Example	`switch (config vlan 10) # ipv6 nd ra hop-limit 70`	
Related Commands		

Notes	

15.1.6.4.12 ipv6 nd ra interval max-period

	ipv6 nd ra interval max-period <time> [min-period <time>] no ipv6 nd ra interval Configures the interval between IPv6 router advertisement (RA) transmissions. The no form of the command resets the parameter to its default value.	
Syntax Description	time	Maximum interval between successive IPv6 router advertisement transmissions Range: 4-1800 seconds
	min-period	Minimum interval between successive IPv6 router advertisement transmissions: • Default is used if no parameter is given • 4-1800
Default	max-period: 600 seconds min-period: See Note	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated syntax description, configuration modes and notes
Example	`switch (config vlan 10) # ipv6 nd ra interval max-period 600`	
Related Commands		
Notes	• The min-period must be 0.33 * <max-period> if <max-period> is >= 9 seconds; otherwise, the default is Router Advertisement Interval • The parameter min-period must be no less than 3 seconds and no greater than 0.75*max-period	

15.1.6.4.13 ipv6 nd ra lifetime

	ipv6 nd ra lifetime <time> no ipv6 nd ra lifetime Router lifetime is associated with a router's usefulness as default route, it does not apply to information contained in other message fields or options. Options that need time limits for their information include their own lifetime fields. The no form of the command resets the parameter to its default value.	
Syntax Description	time	The router lifetime specifies the period that the router can be considered as a default router by RA recipients in seconds. • 0 – the router should not be considered a default router on this interface • 1-9000 – lifetime period advertised in RAs should not be less than the max router advertisement interval
Default	3*<router advertisement interval>	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	

	3.6.4110	Added support for IPv6
Example	`switch (config vlan 10) # ipv6 nd ra lifetime 300`	
Related Commands		
Notes		

15.1.6.4.14 ipv6 nd ra mtu suppress

	ipv6 nd ra mtu suppress no ipv6 nd ra mtu suppress Suppresses advertisement (RA) MTU option sent to router. MTU option ensures all nodes on a link use the same MTU value. The no form of the command restores the MTU option to enabled.	
Syntax Description	N/A	
Default	Suppressed	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated command Syntax and configuration mode
Example	`switch (config vlan 10) # ipv6 nd ra mtu suppress`	
Related Commands		
Notes	If not suppressed, MTU of the interface is advertised.	

15.1.6.4.15 ipv6 nd ra suppress

	ipv6 nd ra suppress [all] no ipv6 nd ra suppress Suppresses periodic and solicited IPv6 router advertisement (RA) transmissions. The no form of the command restores the transmission of RAs.	
Syntax Description	all	Configures the switch to suppress all RAs, including those responding to a router solicitation.
Default	Only unsolicited RAs transmitted periodically are suppressed	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated command syntax and configuration mode
Example	`switch (config vlan 10) # ipv6 nd ra suppress all`	
Related Commands		
Notes		

15.1.6.4.16 ipv6 nd reachable-time

	ipv6 nd reachable-time <time> no ipv6 nd reachable-time Sets the time period the switch includes in the reachable time field of outgoing advertisements (RAs). The no form of the command resets the parameter to its default value.	
Syntax Description	time	In milliseconds; the reachable time defines the period that a node assumes a neighbor is reachable after having received a reachability confirmation. Values: • 0 - unspecified by router • 1 - 3600000 the period that a node assumes a neighbor is reachable
Default	0 (unspecified)	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated command syntax, configuration mode and notes
Example	`switch (config vlan 10) # ipv6 nd reachable-time 30000`	
Related Commands		
Notes	RAs that advertise zero seconds indicate that the router does not specify a reachable time	

15.1.6.4.17 ipv6 nd router-preference

	ipv6 nd router-preference {high | medium | low} no ipv6 nd router-preference Sets the value the switch enters in the default router preference (DRP) field of router advertisements (RAs) it sends. The no form of the command resets the parameter to its default value.	
Syntax Description	N/A	
Default	Medium	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated configuration modes
Example	`switch (config vlan 10) # ipv6 nd router-preference high`	
Related Commands		

Notes	• IPv6 hosts maintain a default router list from which to select a router for traffic to offlink destinations. The router's address is then saved in the destination cache. The neighbor discovery protocol (NDP) prefers routers that are reachable or probably reachable over routers whose reachability is unknown or suspect. For reachable or probably reachable routers, NDP can either select the same router every time or cycle through the router list. DRP values specify a host's preferred router. • If router lifetime is zero, preference value must be medium

15.1.6.4.18 ipv6 nd retrans-timer

	ipv6 nd retrans-timer <time> no ipv6 nd retrans-timer Advertises the time between consecutive neighbor solicitation (NS) messages. The no form of the command resets the parameter to its default value.	
Syntax Description	time	In milliseconds; the time between retransmitted neighbor solicitation messages. Possible values: • 0 – unspecified • Range – 1000-4294967295
Default	0 (unspecified)	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated command syntax, configuration mode and example output
Example	`switch (config vlan 10) # ipv6 nd retrans-timer 1000`	
Related Commands		
Notes		

15.1.6.4.19 ipv6 nd redirects

	ipv6 nd redirects no ipv6 nd redirects Enables sending ICMPv6 redirect messages. The no form of the command disables sending ICMPv6 redirect messages.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface vlan
History	3.4.1100
Example	`switch (config interface vlan 10) # ipv6 nd redirects`
Related Commands	
Notes	

15.1.6.4.20 ipv6 nd dad attempts

	ipv6 nd dad attempts <number> no ipv6 nd dad attempts Sets the number of consecutive neighbor solicitation messages sent for duplicate address detection (DAD) validation. The no form of the command resets the value to its default.	
Syntax Description	number	Number of attempts: • 0 – DAD is not performed • Range: 1-1000
Default	1	
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface	
History	3.4.1100	
	3.6.4110	Updated configuration mode
Role	admin	
Example	`switch (config vlan 10) # ipv6 nd dad attempts 10`	
Related Commands		
Notes		

15.1.6.4.21 clear ipv6 neighbors

	clear ipv6 neighbors {ethernet <slot> /<port> | port-channel <port-channel> | vlan <vlan-id>} [<ipv6-addr>] Removes the specified dynamic IPv6 neighbor discovery cache entries.	
Syntax Description	ethernet	Ethernet port (<slot>/<port>)
	vlan	VLAN interface
	ipv6-addr	IPv6 address
Default	N/A	
Configuration Mode	config	
History	3.4.1100	
	3.6.4110	Updated command
Example	`switch (config) # clear ipv6 neighbors ethernet 1/4`	
Related Commands		
Notes	• Commands that do not specify an IPv6 address remove all dynamic entries for the listed interface • Commands that do not specify an interface remove all dynamic entries	

15.1.6.4.22 ipv6 route

	• General route: ipv6 route [vrf <vrf-name>] {<ipv6-prefix> | <ipv6-address> /<length>} <next-hop-ipv6-address> [<distance>] • Local route: ipv6 route [vrf <vrf-name>] {<ipv6-prefix> | <ipv6-address> /<length>} ipv6 route [<distance>] • Drop route: ipv6 route [vrf <vrf-name>] {<ipv6-prefix> | <ipv6-address> /<length>} null0 [<distance>] • Delete route(s): no ipv6 route [vrf <vrf-name>] {<ipv6-prefix> | <ipv6-address> /<length>} [<next-hop-ipv6-address>] Creates an IPv6 static route. The no form of the command deletes static routes.	
Syntax Description	ipv6-address	IPv6 address
	ipv6-prefix	IPv6 address + mask length without space (e.g. a1:a2::33/64)
	length	Prefix length for the associated address space Range: 1-128
	next-hop-ipv6-address	IPv6 address of the next-hop
	distance	Administrative distance assigned to route. Options include: • No parameter — route is assigned a default administrative distance of 1 • 1-255 — the administrative distance assigned to route
	null0	Creates a black hole route with action DROP
Default	No distance parameter indicated: Administrative distance of 1	
Configuration Mode	config	
History	3.4.1100	
	3.6.4110	Updated command
	3.9.1600	Removed ethernet, port-channel, and vlan options
Example	`switch (config) # ipv6 route 3003:db01:: /64 2001:db01::1`	
Related Commands		
Notes	• Static routes have a default administrative distance of 1 • Assigning a higher administrative distance to a static route configures it to be overridden by dynamic routing data • Multiple routes which are configured to the same destination with the same administrative distance comprise an Equal Cost Multi-Path (ECMP) route • A no command not including a source deletes all statements to the destination • Route with distance value 255 is not inserted to the forwarding table	

15.1.6.4.23 ipv6 routing

	ipv6 routing no ipv6 routing Enables forwarding IPv6 unicast packets. The no form of the command disables IPv6 unicast routing.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.4.1100
Example	`switch (config) # ipv6 routing`
Related Commands	
Notes	When routing is enabled, the switch attempts to deliver inbound packets to destination addresses by forwarding them to interfaces or next hop addresses specified by the IPv6 routing table

15.1.6.4.24 ipv6 routing disable-discard-counter

	ipv6 routing disable-discard-counter no ipv6 routing disable-discard-counter Disables router discard counters. The no form of the command restores discard counters advancing.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.9.2100
Example	`switch (config) # ipv6 routing disable-discard-counter`
Related Commands	
Notes	A reboot is required for the feature to take effect.

15.1.6.4.25 show ipv6 interfaces

	show ipv6 interfaces [{{ethernet <port> | port-channel <port-channel> | vlan <vlan-id>}}| brief] Displays the status of specified routed interfaces that are configured for IPv6.	
Syntax Description	ethernet <port>	Displays output pertaining to the specified Ethernet interface
	port-channel <port-channel>	Displays output pertaining to the specified LAG interface
	vlan <vlan-id>	Displays output pertaining to the specified VLAN interface
	brief	Shows basic IPv6 information regarding all IPv6 interfaces
Default	N/A	
Configuration Mode	Any command mode	

History	3.6.4110
Example	

```
switch (config) # show ipv6 interface

Vlan10 is Enabled ,  line protocol is UP
        IPv6 : Enabled
        Link-local address : fe80::f652:14ff:fe2d:9808
        Global Unicast Addresses :
        2001:db01::2  /64
        Joined Group Addresses :
        ff02::1
        ff02::2
        ff02::1:ff2d:9808
        MTU : 1500 bytes
        ICMP error messages limited to every milliseconds : 100
        ICMP redirects : enabled
        ND DAD : enabled
        Number of DAD attempts : 1
        ND reachable time (milliseconds) : 30000
        ND advertised retransmit interval (milliseconds) : 0
        ND router advertisements maximum interval (seconds) : 600
        ND router advertisements minimum interval (seconds) : 198
        ND router advertisements managed configuration flag : unset
        ND router advertisements other   configuration flag : unset
        ND solicited router advertisement : suppressed
        ND router advertisements lifetime (seconds) : 1800
        ND advertised default router preference : medium
        ND router advertisements hop-limit : 64
```

Related Commands	
Notes	

15.1.6.4.26 show ipv6 interfaces brief

	show ipv6 interfaces [<type> <id>] brief Displays basic IPv6 information regarding all IPv6 interfaces	
Syntax Description	<type> <id>	Specifies the interface for which to display data
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4110	
	3.6.8008	Updated Example

Example

```
 switch (config) # show ipv6 interface brief

------------------------------------------------------------------------------------------------------------
---
Interface   Address/Mask       Primary           Address-state   Admin-state    Oper-state    MTU      VRF
------------------------------------------------------------------------------------------------------------
---
mgmt0       fe80::784e/64                        valid           Enabled        Up            1500
default
Eth1/1      2001::1/64         primary           valid           Enabled        Down          1500
default
Eth1/1      2002::1/64                           valid
```

Related Commands	
Notes	

15.1.6.4.27 show interfaces null0

	show interfaces null0 [vrf <vrf-name>] Displays blackhole route byte and packet counters.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.4110

Example	`switch (config) # show interfaces null0`
	`10 packets`
	`740 bytes`

Related Commands	
Notes	

15.1.6.4.28 show ipv6 neighbors

	show ipv6 neighbors [{ethernet <port> | port-channel <port-channel> | vlan <vlan-id>} | <ipv6 address> | summary] Displays IPv6 neighbor discovery (ND) cache information.	
Syntax Description	ethernet <port>	Displays output pertaining to the specified Ethernet interface.
	vlan <vlan-id>	Displays output pertaining to the specified VLAN interface.
	ipv6 address	IPv6 address of individual neighbor
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.1100	
	3.6.4110	Updated command syntax and Example
Example		

```
switch (config) # show ipv6 neighbors
IPv6 Address               MAC Address        State       Interf
-------------------------  -----------------  ----------  ------
2001:db01::1               f4:52:14:2d:98:88  Reachable   vlan10
```

Related Commands	
Notes	

15.1.6.4.29 show ipv6 route

	show ipv6 route [vrf <vrf-name] {[<ipv6-address> <ipv6-address>/<length> [longer-prefixes]] [connected | bgp | static]} Displays IPv6 neighbor discovery (ND) cache information.	
Syntax Description	ipv6-addr	Filters routes by IPv6 address or prefix
	longer-prefixes	Displays output for longer prefix entries
	connected	Displays entries for routes to networks directly connected to the switch

	static	Displays entries added through CLI commands
	summary	Displays the current contents of the IPv6 routing table in summary format
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.1100	
	3.6.4110	Updated Example
	3.6.8008	Updated Example
Example		

```
switch (config) # show ipv6 route

Flags:
  F: Failed to install in H/W
  B: BFD protected
  i: BFD session initializing
  x: protecting BFD session failed

VRF Name default:
  ------------------------------------------------------------------------------
  Destination        Flag   Gateway           Interface      Source   AD/M
  ------------------------------------------------------------------------------
  fe80::/64                 ::                mgmt0          direct   256/256
  default                   ::              mgmt0          direct   1/1
```

Related Commands	
Note	

15.2 OSPF

Open Shortest Path First (OSPF) is a link-state routing protocol for IP networks. It uses a link state routing algorithm and falls into the group of interior routing protocols, operating within a single autonomous system (AS).

OSPF-speaking routers send Hello packets on all OSPF-enabled IP interfaces. If two routers sharing a common data link agree on certain parameters specified in their respective Hello packets, they become neighbors.

Adjacencies, which can be thought of as virtual point-to-point links, are formed between some neighbors. OSPF defines several network types and several router types. The establishment of an adjacency is determined by the types of routers exchanging Hellos and the type of network over which the Hello packets are exchanged.

Each router sends link-state advertisements (LSAs) over all adjacencies. The LSAs describe all of the router’s links, or interfaces, the router's neighbors, and the state of the links. These links might be to stub networks (those without another router attached), to other OSPF routers, to networks in other areas, or to external networks (those learned from another routing process). Because of the varying types of link-state information, OSPF defines multiple LSA types.

Each router receiving an LSA from a neighbor records the LSA in its link-state database and sends a copy of the LSA to all of its other neighbors. By flooding LSAs throughout an area, all routers will build identical link-state databases.

When the databases are complete, each router uses the SPF algorithm to calculate a loop-free graph describing the shortest (lowest cost) path to every known destination, with itself as the root.

When all link-state information has been flooded to all routers in an area, and neighbors have verified that their databases are identical, it means the link-state databases have been synchronized and the route tables have been built. Hello packets are exchanged between neighbors as keepalives, and LSAs are retransmitted. If the network topology is stable, no other activity should occur.For OSPF network design over NVIDIA L3 VMS, please refer to the Virtual Modular Switch Reference Guide.

15.2.1 Router ID

The router ID is a 32-bit number assigned to the router running the OSPF protocol. This number uniquely identifies the router in the OSPF link-state database.

Router ID can be configured statically, however, if it is not configured, then the default election is as follows:

- If a loopback interface already exists, the router ID selects the highest loopback IP address assigned to a loopback interface. Effective tunnel IP is considered as loopback address.
- Otherwise, the the highest IP address assigned to any other interface on the system is selected as router ID.

15.2.2 ECMP

Equal-cost multi-path (ECMP) routing is a routing strategy where next-hop packet forwarding to a single destination can occur over multiple paths. The OSPF link-state routing algorithm can find multiple routes to the same destination, all multiple routes are added to the routing table only if those routes are equal-cost routes.

In case there are several routes with different costs, only the route with the lowest cost is selected. In case there are multiple routes with the same lowest cost, all of them are used (up to maximum of 64 ECMP routes).

ECMP is not configurable but is enabled by default for OSPF.

15.2.3 Configuring OSPF

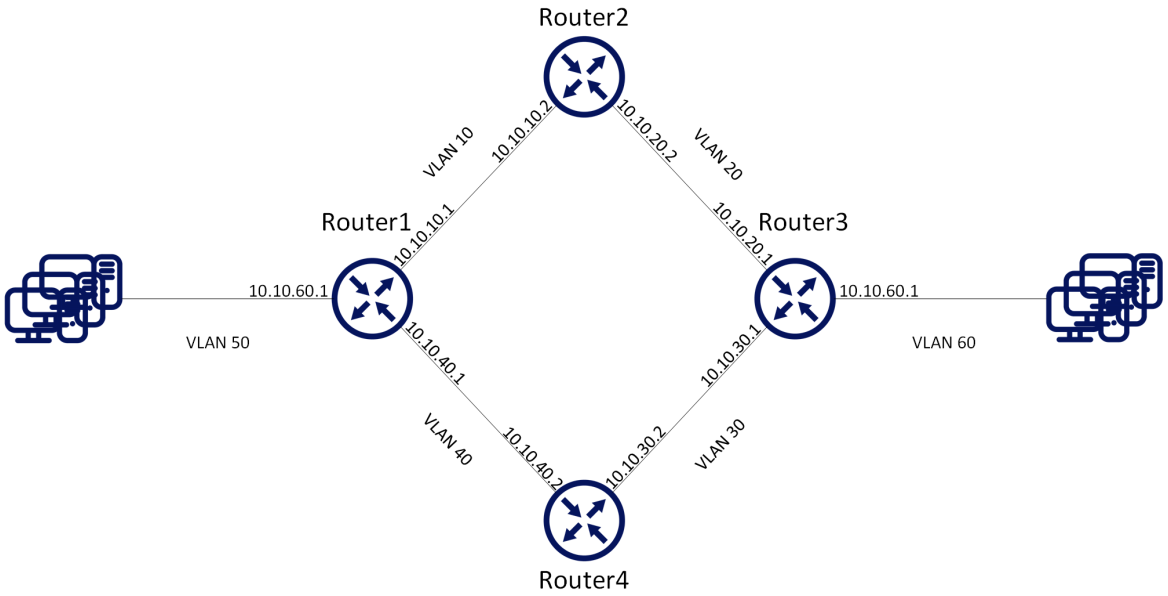

Prerequisites:

> The following configuration example refers to Router 2 in the figure above The remainder of the routers in the figure are configured similarly.

> It is recommended to disable STP before enabling OSPF. Use the command “no spanning-tree”.

1. Enable IP routing functionality. Run:

```
switch (config)# ip routing
```

2. Enable the desired VLAN. Run:

```
switch (config)# vlan 10
switch (config)# vlan 20
```

3. Add this VLAN to the desired interface. Run:

```
switch (config)# interface ethernet 1/1
switch (config ethernet 1/1)# switchport access vlan 10
switch (config ethernet 1/1)# exit
switch (config)# interface ethernet 1/2
switch (config ethernet 1/2)# switchport access vlan 20
```

4. Create a VLAN interface. Run:

```
switch (config)# interface vlan 10
```

5. Apply IP address to the VLAN interface. Run:

```
switch (config interface vlan 10)# ip address 10.10.10.2 /16
```

6. Enable the interface. Run:

```
switch (config interface vlan 10)# no shutdown
```

7. Create a second VLAN interface. Run:

```
switch (config)# interface vlan 20
```

8. Apply IP address to the second VLAN interface. Run:

```
switch (config interface vlan 20)# ip address 10.10.20.2 /16
```

9. Enable the second interface. Run:

```
switch (config interface vlan 20)# no shutdown
```

Basic OSPF Configuration:

1. Enable OSPF configuration commands. Run:

```
switch (config)# protocol ospf
```

2. Create an OSPF instance. Run:

```
switch (config)# router ospf
```

> Only one instance of OSPF per VRF is supported.

3. Associate the VLAN interfaces to the OSPF area. Area 0 is the backbone area. Run:

```
switch (config interface vlan 10)# ip ospf area 0
switch (config interface vlan 10)# exit
switch (config)# interface vlan 20
switch (config interface vlan 20)# ip ospf area 0
```

To verify OSPF configuration and status:

1. Verify OSPF configuration and status. Run:

```
switch (config) # show ip ospf

Routing Process 1 with ID 10.10.10.10 vrf-default

Stateful High Availability disabled
Graceful-restart is not supported
Supports only single TOS (TOS 0) route
Opaque LSA not supported
OSPF Admin State is enabled
Redistributing External Routes: Disabled
Administrative distance 110
Reference Bandwidth is 100Gb
Initial SPF schedule delay 1 msecs
SPF Hold time 10 msecs
Maximum paths to destination 64
Router is not originating router LSA with maximum metric
Condition: Always
Number of external LSAs  0, checksum sum  0
Number of opaque AS LSAs 0,checksum sum 0
Number of areas is 1, 1 normal, 0 stub, 0 nssa
Number of active areas is 1, 1 normal, 0 stub, 0 nssa
```

```
Area (0.0.0.0) (Active)
Interfaces in this area: 2 Active Interfaces: 2
Passive Interfaces: 0
SPF Calculation has run 5 times
This area is Normal area
Number of LSAs: 1, checksum sum 7700
```

2. Verify the OSPF neighbors status. Make sure that each neighbor reaches FULL state with its peer to enable it take part in all dynamic routing changes in the network. Run:

```
switch (config) # show ip ospf neighbors

Neighbor 10.10.10.1, interface address 10.10.10.2
In the area 0.0.0.0 via interface Vlan 10
Neighbor priority is 1, State is FULL
BDR is 10.10.10.1
Options 0
Dead timer due in 35

Neighbor 10.10.20.1, interface address 10.10.20.2
In the area 0.0.0.0 via interface Vlan 20
Neighbor priority is 1, State is FULL
BDR is 10.10.20.1
Options 0
Dead timer due in 35
```

3. Verify the OSPF interface configuration and status. Run:

```
switch (config) # show ip ospf interface

Interface Vlan is 10 Enabled, line protocol is Down
IP address 10.10.10.2, Mask 255.255.0.0 [primary]
Process ID 1 VRF Default, Area 0.0.0.0
OSPF Interface Admin State is enabled
State DOWN, Network Type BROADCAST, Cost 1
Transmit delay 1 sec, Router Priority 1
No designated router on this network
No backup designated router on this network
Timer intervals (sec's): Hello 10, Dead 40, Wait 40, Retransmit 5
No authentication
Number of opaque link LSAs: 0, checksum sum 0

Interface Vlan is 20 Enabled, line protocol is Up
IP address 10.10.20.2, Mask 255.255.0.0 [primary]
Process ID 1 VRF Default, Area 0.0.0.0
OSPF Interface Admin State is enabled
State DESIGNATED ROUTER, Network Type BROADCAST, Cost 1
Transmit delay 1 sec, Router Priority 1
No designated router on this network
No backup designated router on this network
Timer intervals (sec's): Hello 10, Dead 40, Wait 40, Retransmit 5
No authentication
Number of opaque link LSAs: 0, checksum sum 0
```

15.2.4 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community post:

- HowTo Configure OSPF on Switches (Running-Config)

15.2.5 OSPF Commands

- 15.2.5.1 protocol ospf
- 15.2.5.2 router ospf
- 15.2.5.3 router-id
- 15.2.5.4 shutdown
- 15.2.5.5 auto-cost reference-bandwidth
- 15.2.5.6 distance
- 15.2.5.7 redistribute

- 15.2.5.8 timers throttle spf
- 15.2.5.9 area default-cost
- 15.2.5.10 area range
- 15.2.5.11 area stub
- 15.2.5.12 area nssa
- 15.2.5.13 no area
- 15.2.5.14 default-information originate
- 15.2.5.15 summary-address
- 15.2.5.16 ip ospf cost
- 15.2.5.17 ip ospf dead-interval
- 15.2.5.18 ip ospf hello-interval
- 15.2.5.19 ip ospf priority
- 15.2.5.20 ip ospf network
- 15.2.5.21 ip ospf retransmit-interval
- 15.2.5.22 ip ospf passive-interface
- 15.2.5.23 ip ospf transmit-delay
- 15.2.5.24 ip ospf shutdown
- 15.2.5.25 ip ospf authentication
- 15.2.5.26 ip ospf authentication-key
- 15.2.5.27 ip ospf message-digest-key
- 15.2.5.28 ip ospf area
- 15.2.5.29 show ip ospf
- 15.2.5.30 show ip ospf border-routers
- 15.2.5.31 show ip ospf database
- 15.2.5.32 show ip ospf interface
- 15.2.5.33 show ip ospf neighbors
- 15.2.5.34 show ip ospf request-list
- 15.2.5.35 show ip ospf retransmission-list
- 15.2.5.36 show ip ospf summary-address

15.2.5.1 protocol ospf

	protocol ospf no protocol ospf Enables Open Shortest Path First Protocol (OSPF), and unhides the related OSPF commands. The no form of the command deletes the OSPF configuration and hides the OSPF related commands.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.3.3500
Example	`switch (config)# protocol ospf`
Related Commands	ip routing
Notes	

15.2.5.2 router ospf

	router ospf [<process-id> [vrf <vrf-name>]] no router ospf [<process-id> [vrf <vrf-name>]] Creates an ospf instance in the specified VRF and enters the ospf configuration mode. The default process ID is 1 If a VRF is not specified, the OSPF instance is created in the default VRF.	
Syntax Description	process-id	OSPF instance ID
	vrf	VRF name (e.g. default)
Default	Process ID - 1 VRF - active VRF routing-context	
Configuration Mode	config	
History	3.3.3500	
	3.6.1002	Added VRF and process ID parameters and updated Example
Example	`switch (config)# router ospf 2 vrf myvrf` `switch (config router ospf 2)#`	
Related Commands		
Notes	Only one OSPF instance per VRF is supported.	

15.2.5.3 router-id

	router-id <ip-address> no router-id Sets Router ID for the OSPF instance. The no form of the command causes automatic election of router ID by the router.	
Syntax Description	ip-address	The Router ID in IP address format
Default	The router ID is a 32-bit number assigned to the router running the OSPF protocol. This number uniquely identifies the router within an OSPF link-state database. Router ID can be configured statically. However, if it is not configured, then the default election is as follows: • If a loopback interface already exists, the router ID takes the highest loopback IP address assigned to a loopback interface • Otherwise, the highest IP address is elected as router ID	
Configuration Mode	config ospf router	
History	3.3.3500	
	3.7.1100	Updated default
Example	`switch (config router ospf)# router-id 10.10.10.10`	
Related Commands		
Notes		

15.2.5.4 shutdown

	shutdown no shutdown Disables the OSPF instance. The no form of the command enables the OSPF instance.
Syntax Description	N/A
Default	Enable (no shutdown)
Configuration Mode	config ospf router
History	3.3.3500
Example	`switch (config router ospf)# shutdown`
Related Commands	
Note	

15.2.5.5 auto-cost reference-bandwidth

	auto-cost reference-bandwidth <ref-bw> [Gbps | Mbps] no auto-cost reference-bandwidth Configures reference-bandwidth in Gb/s (Default) or Mb/s. The no form of the command resets this parameter to its default value.	
Syntax Description	ref-bw	Range: 1-4294
	Gbps	Value in Gb/s (default if not specified)
	Mbps	Value in Mb/s
Default	100Gbps	
Configuration Mode	config ospf router	
History	3.3.3500	
Example	`switch (config router ospf)# auto-cost reference-bandwidth 10 Gbps`	
Related Commands		
Notes		

15.2.5.6 distance

	distance <value> no distance Configures the OSPF route administrative distance. The no form of the command resets this parameter to default.	
Syntax Description	value	OSPF administrative distance Range is 1-255
Default	110	
Configuration Mode	config ospf router	
History	3.3.3500	
Example	`switch (config router ospf)# distance 100`	

Related Commands	
Notes	

15.2.5.7 redistribute

	redistribute {bgp | direct | static | ebgp | ibgp} no redistribute {bgp | direct | rip | static} Enables importing routes from other routing protocols as well as any statically configured routers into OSPF. The no form of the command disables the importing of the routes.	
Syntax Description	direct	Redistributes directly connected routes
	bgp	Redistributes routes from BGP protocol
	ibgp	Redistributes IBGP routes
	ebgp	Redistributes EBGP routes
	static	Redistributes static configured routes
Default	Disable (no redistribution)	
Configuration Mode	config ospf router	
History	3.6.3506	
Example	`switch (config router ospf)# redistribute direct`	
Related Commands		
Notes	Routes from multiple protocols can be imported in parallel.	

15.2.5.8 timers throttle spf

	timers throttle spf <spf-delay> <spf-hold> no timers throttle spf Sets the OSPF throttle SPF timers. The no form of the command resets the timers to default.	
Syntax Description	spf-delay	The interval by which SPF calculations delayed after a topology change reception Range: 0-100 (milliseconds)
	spf-hold	The minimum delay between two consecutive delay calculations Range: 0-1000 (milliseconds)
Default	spf-delay - 1 millisecond spf-hold - 10 milliseconds	
Configuration Mode	config ospf router	
History	3.3.3500	
Example	`switch (config router ospf)# timers throttle spf 100 1000`	
Related Commands		
Notes		

15.2.5.9 area default-cost

	area <area-id> default-cost <cost> no area <area-id> default-cost Specifies cost for the default summary route sent into an OSPF stub or not-so-stubby area (NSSA). The no form of the command sets the cost to the default value.	
Syntax Description	area-id	OSPF area ID Range: 0-4294967295.
	cost	The cost for the default summary route Range: 1-16777215.
Default	The summary route cost is based on the area border router that generated the summary route	
Configuration Mode	config ospf router	
History	3.3.3500	
Example	`switch (config router ospf)# area 0 default-cost 100`	
Related Commands		
Notes	Base cost for all calculation is 100GbE	

15.2.5.10 area range

	area <area-id> range <ip-address> <prefix> [not-advertise] no area <area-id> range <ip-address> <prefix> [not-advertise] Consolidates and summarizes routes at an OSPF area boundary. The no form of the command removes the ip-prefix range from summarization.	
Syntax Description	area-id	OSPF area ID Range: 0-4294967295
	not-advertise	Suppresses routes that match the specified IP address
	prefix	Network prefix (in the format of /24, or 255.255.255.0 for example)
Default	Disabled	
Configuration Mode	config ospf router	
History	3.3.3500	
Example	`switch (config router ospf)# area 0 range 10.10.10.10 /24`	
Related Commands		
Notes		

15.2.5.11 area stub

	area <area-id> stub [no-summary] no area <area-id> stub [no-summary] Configures an area as an OSPF stub area (an area is created if non-existent). The no form of the command removes the stub area configuration and changes the area to normal, or deletes the area (if stub is not used).

Syntax Description	area-id	OSPF area ID Range: 0-4294967295
	no-summary	Summary route will not be advertised into the stub area
Default	Summary route is advertised	
Configuration Mode	config ospf router	
History	3.3.3500	
Example	`switch (config router ospf)# area 0 stub`	
Related Commands		
Note		

15.2.5.12 area nssa

	area <area-id> nssa [default-information-originate [metric <m-value>] [metric-type <m-type>]] [nosummary] [translate type7 always] no area <area-id> nssa [default-information-originate] [no-summary] [translate type7 always] Configures an area as an OSPF not-so-stubby (NSSA) area. The no form of the command removes the NSSA area configuration and changes the area to default.	
Syntax Description	area-id	OSPF area ID Range: 0-4294967295
	default-information-originate	A default type7 LSA (Link State Advertisements) is generated into the NSSA area
	m-type	Metric type for OSPF Range: 1-2
	m-value	Metric value for OSPF Range: 1-65535
	no-summary	Summary route will not be advertised into the NSSA area
	translate type7 always	Type7 LSAs is translated to type5 LSAs (Link State Advertisements)
Default	Default m-type - 2 Default m-value - 10	
Configuration Mode	config ospf router	
History	3.3.3500	
Example	`switch (config router ospf)# area 0 nssa`	
Related Commands		
Notes	An area can be either stub, NSSA or normal.	

15.2.5.13 no area

	no area <area-id> Deletes OSPF area and its related configuration.

Syntax Description	area-id	OSPF area ID Range: 0-4294967295
Default	N/A	
Configuration Mode	config ospf router	
History	3.3.3500	
Example	`switch (config router ospf)# no area 1`	
Related Commands		
Notes	The command fails if the area is attached to active interfaces	

15.2.5.14 default-information originate

	default-information originate [always] [metric <m-value>] [metric-type <m-type>] no default-information originate Enables default route origination to normal areas. The no form of the command resets the parameter values to their default.	
Syntax Description	always	Default route is always advertised even if the default route is not in the routing table
	metric	Route metric value. Range: 1-65535.
	metric-type	Metric type. Range: 1-2.
Default	m-value – 1 m-type – 2	
Configuration Mode	config ospf router	
History	3.6.8008	
Example	`switch (config router ospf)# default-information originate always`	
Related Commands		
Notes	When default route origination is enabled, the router automatically becomes ASBR and advertises a default route	

15.2.5.15 summary-address

	summary-address <ip-address> <prefix> [not-advertise] no summary-address <ip-address> <prefix> [not-advertise] Creates aggregate addresses for the OSPF protocol. The no form of the command disables the aggregation of the ip-address.	
Syntax Description	ip-address	The summary IP address.
	not-advertise	Suppresses routes that match the specified ip-address.
	prefix	Network prefix (in the format of /24 or 255.255.255.0, for example).
Default	N/A	
Configuration Mode	config ospf router	
History	3.3.3500	
Example	`switch (config router ospf)# summary-address 10.10.10.10 /24`	

Related Commands	
Notes	Maximum of 1500 summarized IP addresses can be configured

15.2.5.16 ip ospf cost

	ip ospf cost <cost> no ip ospf cost Sets OSPF cost of sending packet of this interface. The no form of the command resets this parameter to default.	
Syntax Description	cost	The Interface cost used by the OSPF. Range is 1-65535.
Default	Reference_BW/Link_BW	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.3500	
	3.7.1100	Updated Default
Example	`switch (config interface vlan 10)# ip ospf cost 100`	
Related Commands		
Notes		

15.2.5.17 ip ospf dead-interval

	ip ospf dead-interval <seconds> no ip ospf dead-interval Configures the interval during which at least one Hello packet must be received from a neighbor before the router declares that neighbor as down. The no form of the command resets this parameter to its default.	
Syntax Description	seconds	The dead-interval timer Range: 1-65535 seconds
Default	40 seconds	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.3500	
Example	`switch (config interface vlan 10)# ip ospf dead-interval 10`	
Related Commands		
Notes	The value must be the same for all nodes on the network.	

15.2.5.18 ip ospf hello-interval

	ip ospf hello-interval <seconds> no ip ospf hello-interval Configures the interval between Hello packets that OSPF sends on the interface. The no form of the command resets this parameter to default.

Syntax Description	seconds	The Hello interval timer Range: 1-65535 seconds
Default	10	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.3500	
Example	`switch (config interface vlan 10)# ip ospf hello-interval 20`	
Related Commands		
Notes	The value must be the same for all nodes on the network.	

15.2.5.19 ip ospf priority

	ip ospf priority <number> no ip ospf priority Configures the priority for this OSPF interface. The no form of the command resets this parameter to default.	
Syntax Description	number	The Interface priority used by the OSPF protocol Range: 0-255
Default	1	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.3500	
Example	`switch (config interface vlan 10)# ip ospf priority 100`	
Related Commands		
Notes	• Use the “ip ospf priority” command to set the router priority, which determines the designated router for this network. When two routers are attached to a network, both attempt to become the designated router. • The router with the higher router priority takes precedence. If there is a tie, the router with the higher router ID takes precedence. A router with a router priority set to zero cannot become the designated router or backup designated router.	

15.2.5.20 ip ospf network

	ip ospf network <type> no ip ospf network Sets the OSPF interface network type. The no form of the command resets the interface network type to its default.	
Syntax Description	type	The network type on this interface. • broadcast • point-to-point
Default	Broadcast for VLAN interfaces Point-to-point for router port interfaces	

Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)
History	3.3.3500
Example	`switch (config interface vlan 10)# ip ospf network point-to-point`
Related Commands	
Notes	• The network type influences the behavior of the OSPF interface. An OSPF network type is usually broadcast, which uses OSPF multicast capabilities. Under this network type, a designated router and backup designated router are elected. For point-to-point networks, there are only two neighbors and multicast is not required. • All routers on the same network must have the same network type

15.2.5.21 ip ospf retransmit-interval

	ip ospf retransmit-interval <seconds> no ip ospf retransmit-interval Configures the time between OSPF link-state advertisement (LSA) retransmissions for adjacencies that belongs to the interface. The no form of the command resets this parameter to its default.	
Syntax Description	seconds	The retransmit interval Range: 0-3600 seconds
Default	5	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.3500	
Example	`switch (config interface vlan 10)# ip ospf retransmit-interval 10`	
Related Commands		
Notes		

15.2.5.22 ip ospf passive-interface

	ip ospf passive-interface no ip ospf passive-interface Suppresses flooding of OSPF routing updates on an interface. The no form of the command reverts the status to active OSPF interface.
Syntax Description	N/A
Default	Active interface (no ip ospf passive-interface)
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)
History	3.3.3500
Example	`switch (config interface vlan 10)# ip ospf passive-interface`
Related Commands	

Notes	

15.2.5.23 ip ospf transmit-delay

	ip ospf transmit-delay <seconds> no ip ospf transmit-delay Sets the estimated time required to send an OSPF link-state update packet. The no form of the command resets this parameter to its default.	
Syntax Description	seconds	The transmit-delay interval in seconds Range: 0-3600
Default	1	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.3500	
Example	`switch (config interface vlan 10)# ip ospf transmit-delay 2`	
Related Commands		
Notes		

15.2.5.24 ip ospf shutdown

	ip ospf shutdown no ip ospf shutdown Disables the OSPF instance on the interface. The no form of the command enables the OSPF on this interface.
Syntax Description	N/A
Default	Enabled (no shutdown)
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)
History	3.3.3500
Example	`switch (config interface vlan 10)# ip ospf shutdown`
Related Commands	
Notes	

15.2.5.25 ip ospf authentication

	ip ospf authentication [message-digest] no ip ospf authentication Specifies the authentication type for OSPF. The no form of the command disables the authentication.	
Syntax Description	message-digest	Specifies that message-digest authentication (MD5) is used
Default	Disabled	

Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)
History	3.3.3500
Example	`switch (config interface vlan 10)# ip ospf authentication`
Related Commands	
Notes	• Without message-digest option, a simple password authentication will be used • Message-digest authentication can be enabled only if a key is configured

15.2.5.26 ip ospf authentication-key

	ip ospf authentication-key [<auth-type>] <password> no ip ospf authentication-key To assign a password for simple password authentication for the OSPF. The no form of the command deletes the simple password authentication key.	
Syntax Description	auth-type	The authentication type: • 0 – unencrypted password • 7 – MD5 key
	password	Authentication password (up to 8 alphanumeric string)
Default	Unencrypted password	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.3500	
Example	`switch (config interface vlan 10)# ip ospf authentication-key 0 mycleartextpassword`	
Related Commands		
Notes	• When selecting an encrypted password “7”, the user must input a password encrypted with an MD5 key • When selecting an unencrypted password “0”, the user must input a cleartext password. Then when examining the running-config, it exhibits the encrypted password.	

15.2.5.27 ip ospf message-digest-key

	ip ospf message-digest-key <key-id> md5 [auth-type] <key> no ip ospf message-digest-key <key-id> Sets the message digest key for MD5 authentication. The no form of the command deletes the key for MD5 authentication.	
Syntax Description	auth-type	The authentication type: • 0 – Unencrypted password • 7 – MD5 key
	key	Authentication password, up to 8 alphanumeric string
	key-id	Alphanumeric password of up to 16 bytes
Default	Unencrypted	

Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)
History	3.3.3500
Example	`switch (config interface vlan 10)# ip ospf message-digest-key mykeyid md5 7 mykey`
Related Commands	
Notes	The user cannot delete the last key until authentication is disabled.

15.2.5.28 ip ospf area

	ip ospf area <area-id> no ip ospf area Configures OSPF area of this interface (and creates the area if non-existent). The no form of the command removes the interface from the area.	
Syntax Description	area-id	OSPF area ID Range: 0-4294967295
Default	N/A	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface) config interface loopback	
History	3.3.3500	
Example	`switch (config interface vlan 10)# ip ospf area 0`	
Related Commands		
Notes		

15.2.5.29 show ip ospf

	show ip ospf [<process-id> [vrf <vrf-name>]] Displays general OSPF configuration on specific VRF and status.	
Syntax Description	process-id	OSPF instance ID
	vrf	VRF instance
Default	Process ID - 1 VRF - active VRF routing-context	
Configuration Mode	Any command mode	
History	3.3.3500	
	3.6.1002	Added VRF and process ID parameters and updated Example
Example		

```
switch (config)# show ip ospf 2 vrf myvrf

Routing Process 2 with ID 2.2.2.2 myvrf

Stateful High Availability is not supported
Graceful-restart is not supported
Supports only single TOS (TOS 0) route
Opaque LSA not supported
OSPF Admin State is enabled
Redistributing External Routes: Disabled
Administrative distance 110
Reference Bandwidth is 40 Gbps
Initial SPF schedule delay 1 msecs
SPF Hold time 5000 msecs
Maximum paths to destination 64
Router LSA with maximum metric is not supported
Condition: Always
Number of external LSAs 0, checksum sum 0
Number of opaque AS LSAs 0, checksum sum 0
Number of areas is 1, 1 normal, 0 stub, 0 nssa
Number of active areas is 1, 1 normal, 0 stub, 0 nssa

Area (0.0.0.0) (Active)
Interfaces in this area: 2 Active Interfaces: 2
Passive Interfaces: 0
SPF Calculation has run 6 times
This area is Normal area
Number of LSAs: 3, checksum sum 161346
```

Related Commands	
Notes	

15.2.5.30 show ip ospf border-routers

	show ip ospf border-routers [vrf <vrf-name>] Displays routing table entries to an Area Border Routers.	
Syntax Description	vrf	OSPF routing table entries to an Area Border Routers on specific VRF
Default	VRF - active VRF routing-context	
Configuration Mode	Any command mode	
History	3.3.3500	
	3.6.1002	Added VRF parameter and updated Example

Example

```
switch (config)# show ip ospf border-routers vrf myvrf

OSPF Process ID 2, vrf myvrf Internal Routing Table
Codes: i - Intra-area route, I - Inter-area route
i  1.1.1.1 [0]   ABR   Area: 0.0.0.0,  Next Hop: 21.21.21.1
```

Related Commands	
Notes	

15.2.5.31 show ip ospf database

	show ip ospf database [summary] [<process-id> <area-id> [<link-state-id>]] [adv-router <ip-address> | self-originated] [vrf <vrf-name>] Displays the OSPF database.

Syntax Description	adv-router <ip-address>	Filters per advertise router
	area-id	Filters the command per OSPF area ID Range: 0-4294967295
	link-state-id	The link state ID
	self-originated	Self Originate
	summary	Summarizes the output of the OSPF database
	process-id	Displays OSPF database on specific instance ID
	vrf	Displays OSPF database on specific VRF
Default	Process ID - 1 VRF - active VRF routing-context	
Configuration Mode	Any command mode	
History	3.3.3500	
	3.6.1002	Added VRF and process ID parameters and updated Example
Example	`switch (config)# show ip o`	
Related Commands		

```
switch (config)# show ip ospf database 2 vrf myvrf

OSPF Router with ID (2.2.2.2) (Process ID 2 VRF myvrf)

                Router Link States (Area 0.0.0.0)
                ----------------------------------------
Link ID         ADV Router      Age             Seq             Checksum        LinkCount
2.2.2.2         2.2.2.2         1150            0x80000006      0xbd2a          3
1.1.1.1         1.1.1.1         1152            0x80000006      0xf7f5          3

                Network Link States (Area 0.0.0.0)
                ----------------------------------------
Link ID         ADV Router      Age             Seq             Checksum
21.21.21.2      2.2.2.2         1150            0x80000003      0xbb26
```

Notes	

15.2.5.32 show ip ospf interface

	show ip ospf interface [<process-id>] [vlan <vlan-id> | Ethernet <slot/port | port-channel <number>] [brief] Displays the OSPF related interface configuration.	
Syntax Description	brief	Gives a brief summary of the output
	process-id	Displays OSPF interface configuration on specific instance ID
	vlan <vlan-id>	Displays OSPF interface configuration and status per VLAN interface
	vrf	Displays OSPF interface configuration on specific VRF
Default	Process ID - 1 VRF - active VRF routing-context	
Configuration Mode	Any command mode	
History	3.3.3500	
	3.6.1002	Added VRF and process ID parameters and updated Example

	3.6.4070	Added Ethernet variable

Example

```
switch (config) # show ip ospf interface 2 vrf myvrf

Interface Vlan is 21 Enabled, line protocol is Up
IP address 21.21.21.2,  Mask 255.255.255.0 [primary]
IP address 30.30.30.30, Mask 255.255.255.0
Process ID 2 VRF myvrf, Area 0.0.0.0
OSPF Interface Admin State is enabled
State DESIGNATED ROUTER, Network Type BROADCAST, Cost 10
Transmit delay 1 sec, Router Priority 1
DR is 2.2.2.2
Backup Designated Router is 1.1.1.1
Timer intervals (secs): Hello 10, Dead 40, Wait 40, Retransmit 5
No authentication
Number of opaque link LSAs: 0, checksum sum 0

switch (config) # show ip ospf interface 2 vrf myvrf brief

OSPF Process ID 2 VRF myvrf
Total number of interface: 2
Interface Id    Area            Cost            State           Neighbors       Status
Vlan21          0.0.0.0         10              Enabled         1               Up
Ethernet1/22    0.0.0.0         1               Enabled         1               Up
```

Related Commands	
Notes	

15.2.5.33 show ip ospf neighbors

	show ip ospf [vrf <vrf-name>] neighbors [vlan <vlan-id> | interface <name>] [<neighbor ip address>] Displays the OSPF related interface neighbor configuration.	
Syntax Description	vlan-id	Displays OSPF interface configuration and status per VLAN interface
	neighbor ip address	Filers the output per a specific OSPF neighbor
	vrf	Displays OSPF interface neighbor configuration on specific VRF
Default	VRF - active VRF routing-context	
Configuration Mode	Any command mode	
History	3.3.3500	
	3.6.1002	Added VRF parameter and updated Example
	3.6.4070	Added support for BFD
Example		

```
switch (config) # show ip ospf neighbors vrf myvrf

Neighbor 1.1.1.1, interface address 21.21.21.1
In the area 0.0.0.0  via Interface Vlan 21
Neighbor priority is 1, State is FULL
DR is  2.2.2.2
Backup Designated Router is 1.1.1.1
Options 2
Dead timer due in 36

Neighbor 1.1.1.1, interface address 22.22.22.1
In the area 0.0.0.0  via  1/22
Neighbor priority is 1, State is FULL
No designated router on this network
No backup designated router on this network
Options 2
Dead timer due in 36
switch (config) # show ip ospf neighbors  1/22 vrf myvrf

Neighbor 1.1.1.1, interface address 22.22.22.1
In the area 0.0.0.0  via  1/22
Neighbor priority is 1, State is FULL
No designated router on this network
No backup designated router on this network
Options 2
Dead timer due in 29
```

Related Commands	
Notes	BFD session state is displayed as: established, failed or not established. When BFD is not defined in the command, it is not displayed in the output.

15.2.5.34 show ip ospf request-list

	show ip ospf request-list <neighbor-id> {vlan <vlan-id> | ethernet <slot/port> | port-channel <id>} [vrf <vrf-name>] Displays the OSPF list of all link-state advertisements (LSAs) requested by a router.	
Syntax Description	neighbor-id	Filers the output per a specific OSPF neighbor
	vlan-id	Filers the output per a specific VLAN ID
	vrf <vrf-name>	Displays OSPF request-list on specific VRF
Default	vrf - active VRF routing-context	
Configuration Mode	Any command mode	
History	3.3.3500	

Example

```
switch (config) # show ip ospf request-list 4.4.4.4 vlan 7

OSPF Router with ID (7.7.7.1) (Process ID 1)
Neighbor 4.4.4.4, Interface vlan 7, Address 7.7.7.2
42 LSAs on request-list

Type            LS-ID           ADV-RTR         Seq No          Age             Checksum
1               10.10.10.23     10.10.10.23     0x8000012f      37              0xa7b9
1               10.10.10.24     10.10.10.24     0x8000012f      38              0xbd61
```

Related Commands	
Notes	

15.2.5.35 show ip ospf retransmission-list

	show ip ospf retransmission-list <neighbor-id> {vlan <vlan-id> | ethernet <slot/port> | port-channel <id>} [vrf <vrf-name>] Displays the OSPF list of all link-state advertisements (LSAs) waiting to be resent to neighbors.	
Syntax Description	neighbor-id	Filers the output per a specific OSPF neighbor
	vrf <vrf-name>	Displays OSPF retransmission-list on specific VRF
	vlan-id	Filers the output per a specific VLAN ID
Default	vrf - active VRF routing-context	
Configuration Mode	Any command mode	
History	3.3.3500	
Example		

```
switch (config) # show ip ospf retransmission-list 4.4.4.4 vlan 6

OSPF Router with ID (7.7.7.1) (Process ID 1)
Neighbor 4.4.4.4, Interface vlan 6, Address 6.6.6.2
Link state retransmission due in 3780 msec, Queue length 207

Type            LS-ID           ADV-RTR         Seq No          Age             Checksum
3               22.22.22.22     7.7.7.1         0x80000045      0               0xaaf4
3               192.168.23.2    7.7.7.1         0x80000001      353             0x6752
```

Related Commands	
Notes	

15.2.5.36 show ip ospf summary-address

	show ip ospf summary-address [vrf <vrf-name>] Displays a list of all summary address redistribution information configured on the OSPF.	
Syntax Description	vrf <vrf-name>	Display summary address and area range information on specific VRF
Default	vrf - active VRF routing-context	
Configuration Mode	Any command mode	
History	3.3.3500	
Example		

```
switch (config)# show ip ospf summary-address

OSPF Process ID 1 VRF default
Network         Mask            Area            Advertise       LSA type        Metric          Tag
--------        -----           -----           ----------      ---------       -------         -----
66.66.66.0      255.255.255.0   0.0.0.1         Advertise       Type 3          Auto            N/A
66.66.66.0      255.255.255.0   0.0.0.1         Advertise       Type 7          Auto            N/A
55.55.55.0      255.255.255.0   0.0.0.5         Advertise       Type 3          Auto            N/A
33.33.0.0       255.255.0.0     N/A             Advertise       Type 5          Auto            N/A
44.44.0.0       255.255.0.0     N/A             Advertise       Type 5          Auto            N/A
```

Related Commands	
Notes	

15.3 BGP

Border Gateway Protocol (BGP) is an exterior gateway protocol which is designed to transfer routing information between routers. It maintains and propagates a table of routes which designates network reachability among autonomous systems (ASs).

BGP neighbors, or peers, are routers configured manually to converse using the BGP protocol on top of a TCP session on port 179. A BGP speaker periodically sends keep-alive messages to maintain the connection. Network reachability includes such information as forwarding destinations (IPv4 or IPv6) together with a list of ASs that this information traverses and other attributes, so it becomes possible to construct a graph of AS connectivity without routing loops. BGP makes possible to apply policy rules to enforce connectivity graph.

BGP routers communicate through TCP connection on port 179. Connection between BGP neighbors is configured manually or can be established dynamically by configuring dynamic listen groups. When BGP runs between two peers in the same AS, it is referred to as Internal BGP (iBGP, or Interior Border Gateway Protocol). When it runs between separate ASs, it is called External BGP (eBGP, or Exterior Border Gateway Protocol). Both sides can initiate a connection, after the initial connectivity is created, BGP state machine drives both sides to enter into ESTABLISHED state where they can exchange UPDATE messages with reachability information.

15.3.1 State Machine

In order to make decisions in its operations with peers, a BGP peer uses a simple finite state machine (FSM) that consists of six states: Idle; Connect; Active; OpenSent; OpenConfirm; and Established. For each peer-to-peer session, a BGP implementation maintains a state variable that tracks which of these six states the session is in. The BGP protocol defines the messages that each peer should exchange in order to change the session from one state to another.

The first state is the “Idle” state. In “Idle” state, BGP initializes all resources, refuses all inbound BGP connection attempts and initiates a TCP connection to the peer. The second state is “Connect”. In the “Connect” state, the router awaits the TCP connection to complete and transitions to the “OpenSent” state if successful. If unsuccessful, it initializes the ConnectRetry timer and transitions to the “Active” state upon expiration. In the “Active” state, the router resets the ConnectRetry timer to zero and returns to the “Connect” state. In the “OpenSent” state, the router sends an Open message and waits for one in return in order to transition to the “OpenConfirm” state. KeepAlive messages are exchanged and, upon successful receipt, the router is placed into the “Established” state. In the “Established” state, the router can send/receive: KeepAlive; Update; and Notification messages to/from its peer.

15.3.2 Default Address Family

Default Address Family defines which address family is activated when peer or peer-group becomes active.

When the default address family configuration is modified – it will cause a renegotiation of capabilities for all neighbors that do not have explicit configuration of active address families. The default address family in BGP is IPv4.

15.3.3 Default Route Originate

Default Route Originate initial value is set to “false”.

15.3.4 Peer Groups and Update Groups

Any BGP peer can be defined as part of a peer group and it will inherit peer group configuration or have its own configuration.

A system will automatically generate an update group from peer groups members.

Peer that has a different outbound policy from peer-group will not become a part of update group.

15.3.5 Configuring BGP

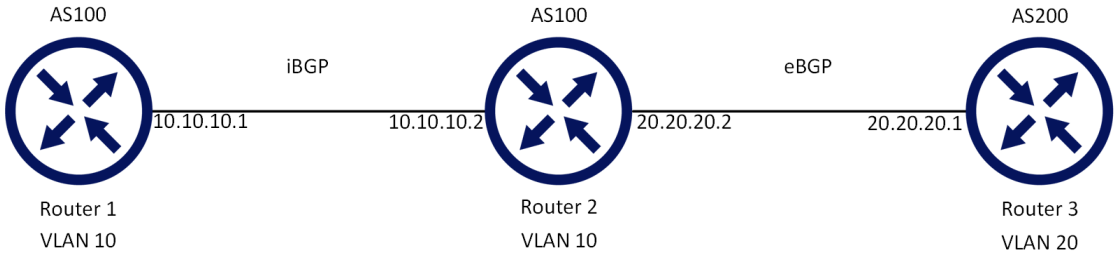

Follow these steps for basic BGP configuration on two switches (Router 1 and Router 2):

Prerequisites:

1. Enable IP routing functionality. Run:

```
switch (config)# ip routing
```

2. Enable the desired VLAN. Run:

```
switch (config)# vlan 10
```

> The same VLAN must be configured on both switches.

3. Add this VLAN to the desired interface. Run:

```
switch (config)# interface ethernet 1/1
switch (config interface ethernet 1/1)# switchport access vlan 10
```

4. Create a VLAN interface. Run:

```
switch (config)# interface vlan 10
```

5. Apply IP address to the VLAN interface on Router 1. Run:

```
switch (config interface vlan 10)# ip address 10.10.10.1 /24
```

6. Apply IP address to the VLAN interface on Router 2. Run:

```
switch (config interface vlan 10)# ip address 10.10.10.2 /24
```

7. Enable the interface. Run:

```
switch (config interface vlan 10)# no shutdown
```

Configure BGP:

1. Enable BGP. Run:

```
switch (config)# protocol bgp
```

2. Configure an AS number that identifies the BGP router. Run:

```
switch (config)# router bgp 100
```

> To run iBGP, the AS number of all remote neighbors should be identical to the local AS number of the configured router.

3. Configure BGP Router 1 neighbor. Run:

```
switch (config router bgp 100)# neighbor 10.10.10.2 remote-as 100
```

4. Configure BGP Router 2 neighbor. Run:

```
switch (config router bgp 100)# neighbor 10.10.10.1 remote-as 100
```

15.3.6 Verifying BGP

1. Check the general status of BGP. Run:

```
switch (config)# show ip bgp summary
BGP router identifier 10.10.10.1, local AS number 100
BGP table version is 100, main routing table version 100
0 network entries using 0 bytes of memory
0 path entries using 0 bytes of memory
0 BGP AS-PATH entries using 0 bytes of memory
0 BGP community entries using 0 bytes of memory
0 BGP extended community entries using 0 bytes of memory
Neighbor        V         AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down     State/PfxRcd
10.10.10.2      0        100     100      76        3    0    0 00:0:10:19 ESTABLISHED
switch (config)#
BGP summary information for VRF default, address family IPv4
```

- Verify that the state of each BGP neighbor reached to ESTABLISHED state.
- If the neighbor is disabled (shutdown). The state of the neighbor will be IDLE.
- BGP incoming and outgoing messages should be incremented.
- The AS number of each neighbor is the correct one.

2. Check the status of the neighbors. Run:

```
switch (config)# show ip bgp neighbors
BGP neighbor is 10.10.10.2, remote AS 100, external link
   BGP version 0, remote router ID 0.0.0.0
   BGP State = ESTABLISHED
   Last read 0:00:00:00, last write 0:00:00:00, hold time is 180, keepalive interval is 60 seconds
   Configured hold time is 180, keepalive interval is 60 seconds
   Minimum holdtime from neighbor is 0 seconds
```

You should be able to see running BGP counters and ESTABLISHED state per active neighbor.

15.3.7 Ethernet Virtual Private Network

Ethernet Virtual Private Network (EVPN) technology provides L2 and L3 VPN services by advertising Ethernet MAC addresses and IP routes over BGP address family. This technology supports multiple forwarding planes including VXLAN.

BGP Layer2-EVPN address family distributes EVPN "routes" between EVPN enabled nodes where some of them are Virtual Tunnel Endpoints (VTEPs) with VXLAN functionality and some of them are transit nodes that perform BGP reflection functionality.

The following route types are defined by RFC 7432:

- MAC/IP advertisement route (route type 2) – advertises MAC and IP addresses of end-systems and their mapping to broadcast domains (VXLAN VNIs and EVPN EVIs). It is used for unicast forwarding, ARP suppression, and advertising default gateway in the EVPN network.
- Inclusive multicast Ethernet tag route (route type 3) – advertises EVPN bridge domain (EVI) and originating router IP address. The EVPN network uses those addresses to instantiate forwarding plane for BUM (Broadcast, unknown Unicast, unknown Multicast) traffic.
- IP prefix route (type 5) – advertises IP prefix, IP gateway, IP address, and HW encapsulation (VNI in the case of VXLAN). This route is used to establish IP prefix LPM routing in the EVPN nodes.

Other route types (type 1 and 4) are used in multi-homing environments only.

RFC 7432 defines BGP attributes that should be used together with Layer-2 EVPN address family routes:

- PMSI tunnel attributes – used for inclusive multicast Ethernet tag route to define multicast type (head end replication) and data path (VNI)
- MAC mobility extended community – used in MAC/IP routes to inform neighbors about MAC roaming events
- Default gateway – used by MAC/IP route to establish default gateway routes
- Route targets – used by all routes to import and export BGP Layer-2 VPN to forwarding and from plane

15.3.8 BGP Unnumbered

BGP unnumbered feature enables a user to establish a BGP session through a P2P Layer-3 link (port or port-channel) without specifying what the IP address of the remote neighbor is, nor what the neighbor's ASN number is.

This Layer-3 link is capable of running IPv6, so the system will use IPv6 link-local addresses that are automatically generated by each IPv6 interface of the local and remote peer. These addresses will be used to establish the BGP TCP session. The ASN number is ignored during the BGP session establishment.

Once IPv6 BGP session is established, the system is able to exchange IPv4 NLRIs (prefixes) over IPv6 BGP session using IPv6 link-local neighbor address as a next hop. The system associates the IPv6 link local address with that neighbor so that the neighbor will be used as a next hop for the routes.

This feature is useful when provisioning a big data center fabric:

- It does not require allocation of an IP subnet on each pair of connected switches
- It simplifies the massive configuration and automation

Remote link-local neighbor address should be available in the local neighbor cache. This address can be populated in any way (ping, static configuration, etc.). It is recommended to use the IPv6 Router Advertisement capability of the router so that the address is populated and refreshed periodically.

Only one neighbor should be available. If more than one exists, one of them is randomly selected.

An ARP entry for 169.254.101.101 is automatically created on each interface on which BGP Unnumbered is configured.

```
switch (config) # show ip arp
VRF Name default:
   Total number of entries: 3
------------------------------------------------------------------------------------
   Address              Type            Hardware Address          Interface
------------------------------------------------------------------------------------
. . .
  169.254.101.101      Static ETH      24:8A:07:7B:85:08         eth 1/17
. . .
```

BGP unnumbered uses 169.254.101.101 as the unnumbered nexthop. As such, while using BGP unnumbered, do not use this address in your topology in the following usages:

1. The interface's IPv4 addresses
2. The prefix or nexthop of static routes
3. The ARP neighbor address

IBGP is not supported for BGP unnumbered.

15.3.9 Configuring BGP Unnumbered

For a basic BGP unnumbered configuration, do the following:

1. Enable IP routing and IPv6 routing

```
ip routing vrf default
ipv6 routing vrf default
```

2. Configure a vrf loopback interface

```
interface loopback 1
interface loopback 1 ip address 25.1.1.1/32 primary
interface vrf default ip address alias loopback1
```

3. Enable IP and IPv6 forwarding on interface

```
interface ethernet 1/2 no switchport force
interface ethernet 1/2 ip enable
interface ethernet 1/2 ipv6 enable
no interface ethernet 1/2 ipv6 nd ra suppress
```

4. Configure BGP

```
protocol bgp
router bgp 200 vrf default
```

5. Enable BGP unnumbered interfaces

```
router bgp 200 vrf default neighbor interface ethernet 1/2
```

6. Test if the session connected well.

```
switch (config) # show ip bgp neighbors interface ethernet 1/2


BGP neighbor: ethernet 1/2 (fe80::268a:7ff:fe7b:8508), remote AS: 100, link: external:
BGP version : 4
Configured hold time in seconds : 180
keepalive interval in seconds (configured) : 60
keepalive interval in seconds (established with peer): 60
Minimum holdtime from neighbor in seconds : 180
Peer group :


Neighbor configuration:
------------------------------------------------------------------------
Configuration IPV4 Unicast IPV6 Unicast L2VPN EVPN
------------------------------------------------------------------------
Configured AFI SAFI Enabled Disabled Disabled
Send Community Disabled Disabled Disabled
Send Extended Community Disabled Disabled Disabled
Route Reflection Disabled Disabled Disabled
Next Hop Unchanged Disabled Disabled Disabled
Extended next hop IPv4 Disabled Enabled Disabled


Neighbor capabilities:
Route Refresh : advertise and received
Enhanced Route Refresh : advertise and received
Soft Reconfiguration : Disabled
Graceful Restart Capability: advertise and received
Address family IPv4 Unicast: advertise and received
Address family IPv6 Unicast: n/a
Address family L2VPN EVPN : n/a
Extended next hop IPv4 : advertise and received


Message statistics:
InQ depth : 0
OutQ depth: 0

......

Connection Information:
Connections established : 1
Dropped : 0
Last Reset : 0:00:00:36
Last Drop Reason : 0 (0)
Maximum hops to external BGP neighbor: 1
Connection State : ESTABLISHED
Local host : fe80::268a:7ff:fe7b:8408
Local port : 43870
Foreign host : ethernet 1/2 (fe80::268a:7ff:fe7b:8508)
Remote port : 179
```

15.3.10 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community posts:

- NVIDIA Onyx BGP Deployment Guide
- How To Configure BGP
- EVPN

15.3.11 BGP Commands

- BGP Commands
- BGP Monitoring Protocol

15.3.12 BGP Commands

- 15.3.12.1 Config
  - 15.3.12.1.1 protocol bgp
  - 15.3.12.1.2 clear ip bgp
  - 15.3.12.1.3 router bgp
- 15.3.12.2 Config Router
  - 15.3.12.2.1 shutdown
  - 15.3.12.2.2 address-family
  - 15.3.12.2.3 aggregate-address
  - 15.3.12.2.4 bestpath as-path multipath-relax
  - 15.3.12.2.5 bgp default
  - 15.3.12.2.6 bgp fast-external-fallover
  - 15.3.12.2.7 bgp listen limit
  - 15.3.12.2.8 bgp listen range peer-group
  - 15.3.12.2.9 bgp redistribute-internal
  - 15.3.12.2.10 cluster-id
  - 15.3.12.2.11 client-to-client reflection
  - 15.3.12.2.12 distance
  - 15.3.12.2.13 graceful-restart stalepath-time
  - 15.3.12.2.14 maximum-paths
  - 15.3.12.2.15 neighbor
  - 15.3.12.2.16 neighbor activate
  - 15.3.12.2.17 neighbor advertisement-interval
  - 15.3.12.2.18 neighbor allowas-in
  - 15.3.12.2.19 neighbor default-originate
  - 15.3.12.2.20 neighbor description
  - 15.3.12.2.21 neighbor ebgp-multihop
  - 15.3.12.2.22 neighbor export-localpref
  - 15.3.12.2.23 neighbor fall-over bfd
  - 15.3.12.2.24 neighbor graceful-restart helper
  - 15.3.12.2.25 neighbor import-localpref
  - 15.3.12.2.26 neighbor local-as
  - 15.3.12.2.27 neighbor maximum-prefix
  - 15.3.12.2.28 neighbor next-hop-peer
  - 15.3.12.2.29 neighbor next-hop-self
  - 15.3.12.2.30 neighbor next-hop-unchanged
  - 15.3.12.2.31 neighbor password
  - 15.3.12.2.32 neighbor no-password
  - 15.3.12.2.33 neighbor peer-group
  - 15.3.12.2.34 neighbor remote-as
  - 15.3.12.2.35 neighbor remove-private-as
  - 15.3.12.2.36 neighbor route-map
  - 15.3.12.2.37 neighbor no-route-map
  - 15.3.12.2.38 neighbor route-reflector-client

        - 15.3.12.2.39 neighbor send-community
        - 15.3.12.2.40 neighbor shutdown
        - 15.3.12.2.41 neighbor soft-reconfiguration
        - 15.3.12.2.42 neighbor soft-reconfiguration inbound
        - 15.3.12.2.43 neighbor timers
        - 15.3.12.2.44 neighbor transport connection-mode passive
        - 15.3.12.2.45 neighbor update-source
        - 15.3.12.2.46 neighbor no-update-source
        - 15.3.12.2.47 neighbor weight
        - 15.3.12.2.48 network
        - 15.3.12.2.49 redistribute
        - 15.3.12.2.50 router-id
        - 15.3.12.2.51 route-map
        - 15.3.12.2.52 timers bgp
        - 15.3.12.2.53 vni
        - 15.3.12.2.54 vni rd
        - 15.3.12.2.55 vni route-target
        - 15.3.12.2.56 vni auto-create
        - 15.3.12.2.57 route-table prefix-list
    - 15.3.12.3 Show
        - 15.3.12.3.1 show {ip | ipv6} bgp
        - 15.3.12.3.2 show ip bgp address-family
        - 15.3.12.3.3 show ip bgp community
        - 15.3.12.3.4 show ip bgp evpn
        - 15.3.12.3.5 show ip bgp evpn summary
        - 15.3.12.3.6 show ip bgp exceptions
        - 15.3.12.3.7 show ip bgp neighbors
        - 15.3.12.3.8 show ip bgp neighbors advertised/received address-family
        - 15.3.12.3.9 show ip bgp neighbors received
        - 15.3.12.3.10 show ip bgp neighbors received detail
        - 15.3.12.3.11 show ip bgp paths
        - 15.3.12.3.12 show ip bgp peer-group
        - 15.3.12.3.13 show ip bgp summary
        - 15.3.12.3.14 show ip bgp update-group
        - 15.3.12.3.15 show ip bgp vrf summary
    - 15.3.12.4 IP AS-Path Access-List
        - 15.3.12.4.1 ip as-path access-list
        - 15.3.12.4.2 show ip as-path access-list
    - 15.3.12.5 IP Community-List
        - 15.3.12.5.1 ip community-list standard
        - 15.3.12.5.2 ip community-list expanded
        - 15.3.12.5.3 show ip community-list

15.3.12.1 Config

15.3.12.1.1 protocol bgp

	protocol bgp no protocol bgp Enables BGPv4, and unhides BGP related commands. The no form of the command deletes all BGP configuration and hides BGP related commands.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.3.5006
Example	`switch (config)# protocol bgp`
Related Commands	ip routing
Notes	

15.3.12.1.2 clear ip bgp

	clear ip bgp <ip-address | ethernet | port-channel | all> [soft] [in | out] Clears BGP learned routes from the BGP table and resets the connection to the neighbor.	
Syntax Description	ip-address	A BGP peer IP address. Only the specified neighbor is reset.
	all	All BGP peers. All BGP neighbors are reset.
	soft	Clears BGP learned routes from the BGP table without resetting the connection to the neighbor
	in	Inbound routes are reset
	out	Outbound routes are reset
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	N/A	
Configuration Mode	config	
History	3.3.5006	First release
	3.3.5200	Updated description
	3.6.3004	Removed “out” parameter
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	`switch (config)# clear ip bgp all` `switch (config)# clear ip bgp vrf default interface ethernet 1/1`	
Related Commands		
Notes	This command removes BGPv4 learned routes from the routing table, reads all routes from designated peers, and sends routes to those peers as required.	

15.3.12.1.3 router bgp

	router bgp <as-number> no router bgp <as-number> Creates and enters a BGP instance with the specified AS number. The no form of the command deletes all router BGP instance configuration.	
Syntax Description	as-number	Autonomous system number: A unique number to be used to identify the AS. The AS is a number which identifies the BGP router to other routers and tags the routing information passed along. Range: 1-4294967295
Default	N/A	
Configuration Mode	config	
History	3.3.5006	
	3.3.5200	Updated syntax description
	3.8.1112	Modified range
Example	`switch (config)# router bgp 100` `switch (config router bgp 100)#`	
Related Commands	ip routing	
Notes		

15.3.12.2 Config Router

15.3.12.2.1 shutdown

	shutdown no shutdown Gracefully disables BGP protocol without removing existing configuration. The no form of the command enables BGP.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config router bgp
History	3.3.5006
Example	`switch (config router bgp 100)# no shutdown`
Related Commands	
Notes	

15.3.12.2.2 address-family

	address-family <ipv4-unicast | ipv6-unicast | l2vpn-evpn> Enables selected address family configuration mode.	
Syntax Description	ipv4-unicast	Enables IPv4 address family configuration mode
	ipv6-unicast	Enables IPv6 address family configuration mode
	l2vpn-evpn	Enables EVPN address family configuration mode

Default	IPv4	
Configuration Mode	config router bgp	
History	3.6.4070	
	3.6.8100	Added “l2vpn-evpn” parameter
Example	`switch (config router bgp 65001) # address-family l2vpn-evpn` `switch (config router bgp 65001 address-family l2vpn-evpn) #`	
Related Commands		
Notes		

15.3.12.2.3 aggregate-address

	aggregate-address <ip_prefix_length> [summary-only] [as-set] [attribute-map] no aggregate-address <ip_prefix_length> [summary-only] [as-set] [attribute-map] Creates an aggregate route in the BGP database. The no form of the command disables ECMP across AS paths.	
Syntax Description	ip_prefix_length	Destination to aggregate
	summary-only	Contributor routes are not advertised
	as-set	Includes AS_PATH information from contributor routes as AS_SET attributes
	attribute-map	Assigns attribute values in set commands of the map’s permit clauses. Deny clauses and match commands in permit clauses are ignored.
Default	Disabled	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.4070	Added support for IPv4 and IPv6
Example	`switch (config router bgp 4) # aggregate-address 3.5.3.7 /32`	
Related Commands		
Notes	• Aggregate routes combine the characteristics of multiple routes into a single route that the switch advertises • Aggregation can reduce the amount of information that a BGP speaker is required to store and transmit when advertising routes to other BGP speakers • Aggregate routes are advertised only after they are redistributed	

15.3.12.2.4 bestpath as-path multipath-relax

	bestpath as-path multipath-relax [force] no bestpath as-path multipath-relax [force] Enables ECMP across AS paths. The no form of the command disables ECMP across AS paths.	
Syntax Description	force	Applies configuration while BGP is admin-up
Default	Disabled	
Configuration Mode	config router bgp	

History	3.3.5006	
	3.3.5200	Updated description and notes
	3.6.3004	Added “force” parameter
Example	`switch (config router bgp 100)# bestpath as-path multipath-relax`	
Related Commands	maximum-paths	
Notes	• With this option disabled, only routes with exactly the same AS path as the best route to a destination are considered for ECMP • With this option enabled, all routes with similar length AS path as the best route are considered for ECMP	

15.3.12.2.5 bgp default

	no bgp default {ipv4-unicast | ipv6-unicast} disable bgp default {ipv4-unicast | ipv6-unicast} Reverts protocol to initial state (IPv4 enabled), enabling setting address families as default for peer or peer-group activation. Disables setting address families as default for peer or peer-group activation.	
Syntax Description	ipv4-unicast	IPv4 unicast address family (enabled by default)
	ipv6-unicast	IPv6 unicast address family (disabled by default)
Default	N/A	
Configuration Mode	config router bgp	
History	3.6.4070	
	3.6.4110	Added support for IPv6
	3.8.1000	Updated command syntax
Example	`switch (config router bgp 100)# bgp default ipv4-unicast`	
Related Commands		
Notes	This command can be used multiple times and each address family can be configured separately.	

15.3.12.2.6 bgp fast-external-fallover

	bgp fast-external-fallover no bgp fast-external-fallover Terminates eBGP sessions of any directly adjacent peer without waiting for the hold-down timer to expire if the link used to reach the peer goes down. The no form of the command waits for hold-down timer to expire before terminating eBGP sessions.
Syntax Description	N/A
Default	no bgp fast-external-fallover
Configuration Mode	config router bgp
History	3.4.0000
Example	`switch (config router bgp 100)# bgp fast-external-fallover`
Related Commands	maximum-paths

Notes	Although this feature improves BGP conversion time, it may cause instability in your BGP table due to a flapping interface.

15.3.12.2.7 bgp listen limit

	bgp listen limit <maximum> no bgp listen limit Limits the number of dynamic BGP peers allowed on the switch. The no form of the command resets to the default value.	
Syntax Description	maximum	The maximum number of dynamic BGP peers to be allowed on the switch Range: 1-128
Default	100	
Configuration Mode	config router bgp	
History	3.4.0000	
Example	`switch (config router bgp 100)# bgp listen limit 101`	
Related Commands		
Notes		

15.3.12.2.8 bgp listen range peer-group

	bgp listen range <ip-prefix> peer-group <peer-group-name> remote-as <as-number> no bgp listen range <ip-prefix> <length> Identifies a range of IP addresses from which the switch will accept incoming dynamic BGP peering requests. After applying the no form of the command, the switch will no longer accept dynamic peering requests on the range.	
Syntax Description	ip-address	IP address
	length	Mask length (e.g. /24 or 255.255.255.254)
	peer-group-name	Peer group name
	remote-as <as-number>	Remote peer's number
Default	100	
Configuration Mode	config router bgp	
History	3.4.0000	
Example	`switch (config router bgp 100)# bgp listen range 10.10.10.10 /24 peer-group my-group remote-as 13`	
Related Commands		

Notes	• To create a static peer group, use the command *neighbor peer-group* • Neighbors in a dynamic peer group are configured as a group and cannot be configured individually • The no form of the command may take up to a few seconds to take effect if there are many dynamic peers and/or a lot of routes. While the clean-up process is running, creation of a new listen range that overlaps the deleted one will fail. • If dynamic peer range is defined with an overlap to another defined range, the longest remote address prefix take affect

15.3.12.2.9 bgp redistribute-internal

	bgp redistribute-internal no bgp redistribute-internal Enables iBGP redistribution into an interior gateway protocol (IGP). The no form of the command disables iBGP redistribution into an interior gateway protocol (IGP).	
Syntax Description	ip-prefix	IP address
	length	Mask length (e.g. /24 or 255.255.255.254)
	peer-group-name	Peer group name
	remote-as <as-number>	Remote peer's number
Default	Disabled	
Configuration Mode	config router bgp	
History	3.4.0000	
Example	`switch (config router bgp 100)# bgp redistribute-internal`	
Related Commands		
Notes		

15.3.12.2.10 cluster-id

	cluster-id <ip-address> [force] no cluster-id <ip-address> [force] Configures the cluster ID in a cluster with multiple route reflectors. The no form of the command resets the cluster ID for route reflector.	
Syntax Description	ip-address	The route reflector cluster ID. • 0.0.0.1 to 255.255.255.255 Valid cluster ID number • 0.0.0.0 removes the cluster-ID from the switch (similar to "no cluster-id")
	force	Applies configuration while BGP is admin-up
Default	Cluster ID is the same as Router ID	
Configuration Mode	config router bgp	
History	3.2.1000	
	3.4.0000	Updated syntax description
	3.6.3004	Added "force" parameter

Example	`switch (config router bgp 100)# cluster-id 10.10.10.10`
Related Commands	
Notes	

15.3.12.2.11 client-to-client reflection

	client-to-client reflection no client-to-client reflection The switch will be configured as a route reflector. The no form of the command stops the switch from being a route reflector
Syntax Description	N/A
Default	client-to-client reflection is enabled
Configuration Mode	config router bgp
History	3.2.1000
Example	`switch (config router bgp 100)# client-to-client reflection`
Related Commands	
Notes	

15.3.12.2.12 distance

	distance <external> <internal> <local> no distance Sets the administrative distance of the routes learned through BGP. The no form of the command resets the administrative distance its default.	
Syntax Description	external	Administrative distance for external BGP routes Range: 1-255
	internal	Administrative distance for internal BGP routes Range: 1-255
	local	Administrative distance for local BGP routes Range: 1-255
Default	external: 20 internal: 200 local: 200	
Configuration Mode	config router bgp	
History	3.3.5006	
Example	`switch (config router bgp 100)# distance 10 20 30`	
Related Commands		

Notes	• Routers use administrative distances to decide on a route when two protocols provide routing information to the same destination • Lower distance values correspond to higher reliability • Routes are external when learned from an external autonomous system • Routes are internal when learned from a peer in the local autonomous system • Local routes are those networks listed with a network router configuration command, often as back doors, for the router or for the networks being redistributed from another process • BGP routing tables do not include routes with a distance of 255

15.3.12.2.13 graceful-restart stalepath-time

	graceful-restart stalepath-time <interval> no graceful-restart stalepath-time Configures the maximum time that stale routes from a restarting BGP neighbor are retained after a BGP session is reestablished with that peer. The no form of the command resets to the default value.	
Syntax Description	interval	Time in seconds Range: 1-3600
Default	300 seconds	
Configuration Mode	config router bgp	
History	3.4.0000	
Example	`switch (config router bgp 100)# graceful-restart stalepath-time 350`	
Related Commands		
Note		

15.3.12.2.14 maximum-paths

	maximum-paths [ibgp] <maximum-path> Configures the maximum number of parallel eBGP/iBGP routes that the switch installs in the routing table.	
Syntax Description	ibgp	Sets the configuration on the internal BGP
	maximum-path	The number of routes to install to the routing table Range: 1-32
Default	1	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.3.5200	Updated description and notes
	3.6.4070	Updated maximum-path range
Example	`switch (config router bgp 100)# maximum-paths ibgp 10`	
Related Commands		

Notes	• This command provides an ECMP parameter that controls the number of equal-cost paths that the switch installs in the routing table for each destination • The action is effective after BGP restart • If the parameter “ibgp” is not used, the setting is applied on routes learned from peers from other ASs • If “ibgp” is used, the setting is applied to routes learned from peers of the same AS

15.3.12.2.15 neighbor

	neighbor <ethernet | port-channel> no neighbor <ethernet | port-channel> Configures a neighbor. The no form of the command removes the neighbor, dropping the connection and all routes if already connected.	
Syntax Description	ethernet	Ethernet type interface
	port-channel	LAG type interface
Default	Disabled	
Configuration Mode	config config router bgp	
History	3.9.0500	
Example	`switch (config router bgp 100)# neighbor interface ethernet 1/17`	
Related Commands	ip routing router bgp	
Notes	• This command supports BGP unnumbered neighbors • IBGP is not supported. For incoming IBGP connection request, it will be rejected and a warning will be logged	

15.3.12.2.16 neighbor activate

	neighbor <ip-address | peer-group | ethernet | port-channel> activate no neighbor <ip-address | peer-group | ethernet | port-channel> activate disable neighbor <ip-address | peer-group | ethernet | port-channel> activate Sends advertisement for given address-family to neighbor. The no form of the command removes the command from running-config and enables inheritance. The disable form of the command sets boolean value to false and disables inheritance.	
Syntax Description	ip-address	Neighbor IP address
	peer-group	Peer group name
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	N/A	
Configuration Mode	config router bgp config router bgp address-family	
History	3.6.4070	
	3.6.4110	Added “disable” option to the command

	3.6.8100	Added “config router bgp address-family” configuration mode
	3.9.0300	Added support for unnumbered neighbors and Updated example

Example	`switch (config router bgp 100)# no neighbor 10.10.10.1 activate`
	`switch (config router bgp 65001 address-family l2vpn-evpn)# neighbor 192.168.3.2 activate`
	`switch (config router bgp 200)# vrf default address-family ipv4-unicast neighbor interface ethernet 1/1 activate`	

Related Commands		

Notes	There are 4 possible ways of using the “disable” prefix:
	• At the beginning of the command
	`switch (config) # disable router bgp 65001 address-family l2vpn-evpn neighbor 192.168.3.2 activate`
	• At the end of the command
	`switch (config) # router bgp 65001 address-family l2vpn-evpn neighbor 192.168.3.2 activate disable`
	• After the “router bgp *”
	`switch (config) # router bgp 65001 disable address-family l2vpn-evpn neighbor 192.168.3.2 activate`
	• After the “router bgp * address-family l2vpn-evpn”
	`switch (config) # router bgp 65001 address-family l2vpn-evpn disable neighbor 192.168.3.2 activate`	

15.3.12.2.17 neighbor advertisement-interval

	neighbor <ip-address | peer-group-name | ethernet | port-channel> advertisement-interval <delay> no neighbor <ip-address | peer-group-name | ethernet | port-channel> advertisement-interval Sets the minimum route advertisement interval (MRAI) between the sending of BGP routing updates. The no form of the command disables this function.	
Syntax Description	ipv4_addr, ipv6_addr	A BGP peer IP address
	peer-group-name	Peer group name
	delay	Time (in seconds) is specified by an integer Range: 0-600; where “0” disables this function and prevents the system from inheriting this parameter’s group configuration
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	30 seconds	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.3004	Updated description of “delay” parameter
	3.9.0300	Added support for unnumbered neighbors and Updated example

Example	`switch (config router bgp 100)# neighbor 10.10.10.10 advertisement-interval 100`
	Without address family:
	`switch (config router bgp 200)# vrf default neighbor interface ethernet 1/3 advertisement-interval 7`
	With address family—can be done only on peer group not on single neighbor:
	`switch (config router bgp 200)# vrf default address-family ipv4-unicast neighbor interface ethernet 1/3 advertisement-interval 7`

Related Commands	
Notes	When configuring an advertisement interval to a BGP session, this interval is implemented per prefix route of that session. For example: If a session is configured with advertisement interval of 100 seconds, when it first learns a new route it automatically sends an update on this route. If it learns another route in the same prefix as the initial route, it waits for 100 seconds. But if it learns another route in a different prefix it immediately advertises that route and does not wait another 100 seconds.

15.3.12.2.18 neighbor allowas-in

	neighbor <ip-address | peer-group-name | ethernet | port-channel> allowas-in [number] no neighbor <ip-address | peer-group-name | ethernet | port-channel> allowas-in Configures the switch to permit the advertisement of prefixes containing duplicate autonomous switch numbers (ASNs). The no form of the command disables this function.	
Syntax Description	ip-address	A BGP peer IP address
	peer-group-name	Peer group name
	number	Number of switch's (ASN) allowed in path Range: 0-10; where "0" disables this function and prevents the system from inheriting this parameter's group configuration
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	N/A	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.3004	Updated description of "number" parameter
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 allowas-in 2` `switch (config router bgp 200)# vrf default neighbor interface ethernet 1/1 allowas-in`	
Related Commands	ip routing router bgp <as-number>	
Notes	Neighbors from the same AS as the router are considered as iBGP peers, and neighbors from other ASs are considered eBGP peers.	

15.3.12.2.19 neighbor default-originate

	neighbor <ip-address | peer-group | ethernet | port-channel> default-originate [route_map_name] no neighbor <ip-address | peer-group | ethernet | port-channel> default-originate [route_map_name] disable neighbor <ip-address | peer-group | ethernet | port-channel> default-originate [route_map_name] Enables advertisement of the default route to a specified neighbor or peer group. The no form of the command disables advertisement of the default route and enables inheritance. The disable form of the command disables advertisement of the default route and disables inheritance.	
Syntax Description	ip-address	Neighbor IPv4 address
	peer-group	Peer group name
	route_map_name	Route map name that modifies default route attributes
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	N/A	
Configuration Mode	config router bgp	
History	3.6.4070	
	3.6.4110	Added “disable” option to the command
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	`switch (config router bgp 100)# neighbor 10.10.10.1 default-originate default-attr` `switch (config router bgp 200)# vrf default address-family ipv4-unicast neighbor interface ethernet 1/1 default-originate`	
Related Commands		
Notes		

15.3.12.2.20 neighbor description

	neighbor <ip-address | peer-group-name | ethernet | port-channel> description <string> no neighbor <ip-address | peer-group-name | ethernet | port-channel> description Associates descriptive text with the specified peer or peer group. The no form of the command removes the description from the peer.	
Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	string	Free string, up to 80 characters in length
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	No description	

Configuration Mode	config router bgp	
History	3.3.5006	
	3.6.4070	Added support for IPv6 and IPv4
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 description The next door neighbor` `switch (config router bgp 200)# vrf default neighbor interface ethernet 1/1 description test desc`	
Related Commands		
Notes	The peer description only appears in the show commands	

15.3.12.2.21 neighbor ebgp-multihop

	neighbor <ip-address | peer-group-name> ebgp-multihop [<ttl>] no neighbor <ip-address | peer-group-name> ebgp-multihop Enables BGP to connect to external peers that are not directly connected to the switch. The no form of the command resets the value to the default (TTL = 1).	
Syntax Description	ip-address	IP address of the BGP-speaking neighbor
	peer-group-name	Peer group name
	ttl	Time-to-live Range: 1-255 hops; where “1” disables connecting to external peers and prevents the system from inheriting this parameter’s group configuration
Default	ttl—1	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.3.5200	Updated Default
	3.6.3004	Updated description of “ttl” parameter
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 ebgp-multihop 5`	
Related Commands	ip routing neighbor <ip-address> remote-as <as-number>	
Notes	The command does not establish the multi-hop if the only route to the peer is the default route (0.0.0.0)	

15.3.12.2.22 neighbor export-localpref

	neighbor <ip-address | peer-group-name | ethernet | port-channel> export-localpref <value> no neighbor <ip-address | peer-group-name | ethernet | port-channel> export-localpref Configures the local preference value sent to the specified peer or peer group. The no form of the command resets the local preference to its default value.	
Syntax Description	ip-address	IP address of the BGP-speaking neighbor

	peer-group-name	Peer group name
	value	Preference value Range: 0-2147483647; where “100” configures the default, and prevents the system from inheriting this parameter’s group configuration
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	100	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.3004	Updated description of “value” parameter
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 export-localpref 100` `switch (config router bgp 200)# vrf default neighbor interface ethernet 1/1 export-localpref 66`	
Related Commands		
Notes		

15.3.12.2.23 neighbor fall-over bfd

	neighbor <ip-address| ip-address | peer-group-name> fall-over bfd no neighbor <ip-address | ip-address | peer-group-name> fall-over bfd Disables BFD as a mechanism to detect failure. The no form of the command enables BFD neighbor.	
Syntax Description	peer-group-name	Peer group name
	ip-address	IP address of the neighbor
Default	Enabled	
Configuration Mode	config router bgp	
History	3.6.4070	
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 bfd`	
Related Commands		
Notes	The command “no neighbor <ip_address> fall-over bfd” affects traffic. BGP will restore the connection based on Hello protocol.	

15.3.12.2.24 neighbor graceful-restart helper

	neighbor <ip-address | peer-group-name> graceful-restart helper no neighbor <ip-address | peer-group-name> graceful-restart helper Enables BGP graceful restart helper mode for the specified BGP neighbor or peer group. The no form of the command disables this parameter.	
Syntax Description	ip-address	IP address of the BGP-speaking neighbor
	peer-group-name	Peer group name

Default	Graceful restart is enabled
Configuration Mode	config router bgp
History	3.4.0000
Example	`switch (config router bgp 100)# neighbor graceful-restart helper`
Related Commands	
Notes	• When graceful restart helper mode is enabled, the switch retains routes from neighbors capable of graceful restart while those neighbors are restarting BGP • Individual neighbor configuration takes precedence over the global configuration

15.3.12.2.25 neighbor import-localpref

	neighbor <ip-address | peer-group-name | ethernet | port-channel> import-localpref <value> no neighbor <ip-address | peer-group-name | ethernet | port-channel> import-localpref <value> Configures the local preference value assigned to routes received from the specified peer or peer group. The no form of the command resets the local preference to its default value.	
Syntax Description	ip-address	IP address of the BGP-speaking neighbor
	peer-group-name	Peer group name
	value	Preference value Range: 0-2147483647; where “100” configures the default, and prevents the system from inheriting this parameter’s group configuration
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	100	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.3004	Updated description of “value” parameter
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 import-localpref 100` `switch (config router bgp 200)# vrf default neighbor interface ethernet 1/1 import-localpref 55`	
Related Commands		
Notes		

15.3.12.2.26 neighbor local-as

	neighbor <ip-address | peer-group-name | ethernet | port-channel> local-as <asn-id> [no-prepend | no-prepend replace-as] no neighbor <ip-address | peer-group-name | ethernet | port-channel> local-as Enables the modification of the AS path attribute for routes received from an eBGP neighbor. The no form of the command disables AS path modification for the specified peer or peer group.	
Syntax Description	ip-address	IP address of the BGP-speaking neighbor
	peer-group-name	Peer group name
	asn-id	AS number that is sent instead of the actual AS of the switch. Range: 0-4294967295
	no-prepend	local-as number is not prepended to the routes received from external neighbors
	no-prepend replace-as	Replaces the local-as (as configured with the IP address argument) in the AS path attribute without pre-pending it to the routes received from external neighbors.
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	N/A	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.3004	Updated description of “as-id” parameter
	3.6.4070	Added support for IPv6 and IPv4
	3.6.4110	Updated command syntax
	3.8.2000	Modified the "replace-as" option and changed it to "no-prepend replace-as"
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	`switch (config router bgp 4)# neighbor 100.100.100.100 local-as 123` `switch (config router bgp 200)# vrf default address-family ipv4-unicast neighbor interface ethernet 1/1 send-community`	
Related Commands	ip routing neighbor <ip-address> remote-as <as-number>	
Notes	• This function allows the switch to appear as a member of a different autonomous system (AS) to external peers • To disable peering with the neighbor run the command “clear ip bgp”	

15.3.12.2.27 neighbor maximum-prefix

	neighbor <ip-address | peer-group-name | ethernet | port-channel> maximum-prefix <maximum> [warning-only] no neighbor <ip-address | peer-group-name | ethernet | port-channel> maximum-prefix Configures the number of BGP routes the switch accepts from a specified neighbor and defines an action when the limit is exceeded. The no form of the command removes the limitation.	
Syntax Description	ip-address	IP address of the BGP-speaking neighbor
	peer-group-name	Peer group name
	maximum	Number of BGP routes the switch accepts from a specified neighbor Range: 1-2147483647; where “12000” configures the default, and prevents the system from inheriting this parameter’s group configuration
	warning-only	Only generates a warning rather than disconnecting the neighbor
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	12000	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.3004	Updated description of “maximum” parameter
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 maximum-prefix 12000 warning-only` `switch (config router bgp 200)# vrf default  neighbor interface ethernet 1/1 maximum-prefix 88`	
Related Commands	ip routing neighbor <ip-address> remote-as <as-number>	
Notes		

15.3.12.2.28 neighbor next-hop-peer

	neighbor <ip-address | peer-group-name | ethernet | port-channel> next-hop-peer [disable] no neighbor <ip-address | peer-group-name | ethernet | port-channel> next-hop-peer Configures the switch to replace the next-hop attribute in routes advertised to IBGP peers with the address of the EBGP peer that advertised this route. The no form of the command disables this function.	
Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name

	disable	Disables this function and prevents the system from inheriting this parameter's group configuration
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	no next-hop-peer	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.6.3004	Added "disable" parameter
	3.9.0300	Added support for unnumbered neighbors, updated command description, and Updated example
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 next-hop-peer` `switch (config router bgp 200)# vrf default neighbor interface ethernet 1/1 next-hop-peer`	
Related Commands		
Notes	This command overrides the next hop for all routes received from this neighbor or peer group	

15.3.12.2.29 neighbor next-hop-self

	neighbor <ip-address | peer-group-name | ethernet | port-channel> next-hop-self [disable] no neighbor <ip-address | peer-group-name | ethernet | port-channel> next-hop-self Configures the IP address of the router as the next hop address in routes advertises to the specific neighbor. The no form of the command resets this parameter to its default.	
Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	disable	Disables this function and prevents the system from inheriting this parameter's group configuration
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	no next-hop-self	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.6.4070	Added support for IPv6
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 next-hop-self` `switch (config router bgp 200)# vrf default neighbor interface ethernet 1/1 next-hop-self`	
Related Commands	neighbor <ip-address> remote-as <as-number>	

Notes	• This function is used in networks where BGP neighbors do not directly access all other neighbors on the same subnet. • In the default state, the next hop is generated based on the IP address and the present next hop in the route information.

15.3.12.2.30 neighbor next-hop-unchanged

	neighbor <ip-address | peer group | ethernet | port-channel> next-hop-unchanged no neighbor <ip-address | peer group | ethernet | port-channel> next-hop-unchanged disable neighbor <ip-address | peer group | ethernet | port-channel> next-hop-unchanged Enables preserving BGP next-hop when forwarding routes to this eBGP peer or all eBGP peers in this address family. The no form of the command removes configuration and enables inheritance of AFI SAFI next-hop-unchanged configuration from a peer group if this neighbor is member in one. The disable form of the command disables preserving BGP next-hop when forwarding routes to this eBGP peer or all eBGP peers in this address family.	
Syntax Description	ip-address	Neighbor IP address
	peer_group	Peer group name
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	The next-hop of a route is preserved when advertising the route to an iBGP peer, but is updated when advertising the route to an eBGP peer. Setting this to “true” overrides this behavior and preserves the next-hop when routes are advertised to this eBGP peer.	
Configuration Mode	config router bgp address-family	
History	3.6.8100	
	3.9.0300	Added support for unnumbered neighbors and Updated example

Example	`switch (config router bgp 65001 address-family l2vpn-evpn) # neighbor 192.168.5.2 next-hop-unchanged`
	`switch (config router bgp 65001 address-family l2vpn-evpn) # next-hop-unchanged`
	`switch (config router bgp 200)# vrf default address-family ipv4-unicast neighbor interface ethernet 1/1 next-hop-unchanged`	

Related Commands	address-family l2vpn-evpn	
Note	There are 4 possible ways of using the “disable” prefix: • At the beginning of the command switch (config) # disable router bgp 65001 address-family l2vpn-evpn neighbor 192.168.3.2 next-hop-unchanged • At the end of the command switch (config) # router bgp 65001 address-family l2vpn-evpn neighbor 192.168.3.2 next-hop-unchanged disable • After the “router bgp *” switch (config) # router bgp 65001 disable address-family l2vpn-evpn neighbor 192.168.3.2 next-hop-unchanged • After the “router bgp * address-family l2vpn-evpn” switch (config) # router bgp 65001 address-family l2vpn-evpn disable neighbor 192.168.3.2 next-hop-unchanged	

15.3.12.2.31 neighbor password

	neighbor <ip-address | peer-group-name | ethernet | port-channel> password [<encryption>] <string> no neighbor <ip-address | peer-group-name | ethernet | port-channel> password Enables authentication on a TCP connection with a BGP peer. The no form of the command resets the value to its default.	
Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	encryption	Possible values: • no parameter - clear text • 0—clear text • 7—obfuscated
	string	Up to 8 bytes in length
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	no neighbor password	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 password 7 admin123` `switch (config router bgp 200)# vrf default neighbor interface ethernet 1/1 password 0 test`	
Related Commands		
Note	• Peers must use the same password to ensure communication • “neighbor <ip-address> password 7 <password>” can only accept data that was created using “show config” • “show config” will never show the clear-test password, it will always be obfuscated (and thus displayed using the 'password 7' syntax). • Router BGP neighbor password cannot be set when enabling secure mode • Router BGP peer-group password cannot be set when enabling with secure mode	

15.3.12.2.32 neighbor no-password

	neighbor <ip-address | peer-group-name | ethernet | port-channel> no-password Disables authentication for peer without inheritance.	
Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	N/A	
Configuration Mode	config router bgp	
History	3.6.3004	

	3.9.0300	Added support for unnumbered neighbors
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 no-password`	
Related Commands	neighbor password	
Notes		

15.3.12.2.33 neighbor peer-group

	1. neighbor <ip-address | ethernet | port-channel> peer-group <peer-group-name> 2. neighbor <peer-group-name> peer-group 3. no neighbor <ip-address | ethernet | port-channel> peer-group <peer-group-name> 4. no neighbor <peer-group-name> peer-group 1. Assigns BGP neighbors to an existing peer group 2. Creates a peer-group 3. Unassigns a BGP neighbor from a peer-group 4. Deletes the peer-group	
Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	N/A	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.3004	Added notes
	3.6.4070	Added support for IPv6 and IPv4
	3.9.0300	Added support for unnumbered neighbors and modified note
Example	`switch (config router bgp 100)# neighbor groupA peer-group` `switch (config router bgp 100)# neighbor 1.2.3.4 peer-group groupA`	
Related Commands		
Notes	• Once a peer group is created, the group name can be used as a parameter in neighbor configuration commands, and the configuration will be applied to all members of the group • Settings applied to an individual neighbor in the peer group override group settings • A neighbor can only belong to one peer group, so issuing this command for a neighbor that is already a member of another group removes it from that group • When a neighbor is removed from a peer group, the neighbor does not retain the configuration inherited from the peer group. • Router BGP peer-group password cannot be set when enabling with secure mode • A BGP group must be used by either a single listen range, or by a set of neighbors sharing the same type (iBGP or eBGP) • A group must already exist before a node is configured to use it • Any configuration change on a group affects each of the peers inheriting this specific parameter from the group only after undergoing admin state toggle	

15.3.12.2.34 neighbor remote-as

	neighbor <ip-address> remote-as <as-number> no neighbor <ip-address> remote-as <as-number> Configures a neighbor. The no form of the command removes the neighbor, dropping the connection and all routes if already connected.	
Syntax Description	ipv4_addr, ipv6_addr	IP address of the neighbor
	as-number	The BGP peer as-number Range: 1-65535
Default	N/A	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.3.5200	Updated description and note
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 remote-as 200`	
Related Commands	ip routing router bgp <as-number>	
Notes	Neighbors from the same AS as the router are considered as iBGP peers, and neighbors from other ASs are considered eBGP peers	

15.3.12.2.35 neighbor remove-private-as

	neighbor <ip-address | peer-group-name | ethernet | port-channel> remove-private-as [disable] no neighbor <ip-address | peer-group-name | ethernet | port-channel> remove-private-as Removes private autonomous system numbers from outbound routing updates for external BGP (eBGP) neighbors. The no form of the command preserves private AS numbers for the specified peer.	
Syntax Description	ipv4_addr, ipv6_addr	A BGP peer IP address
	peer-group-name	Peer group name
	disable	Preserves private AS numbers for the specified peer and prevents the system from inheriting this parameter's group configuration
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	N/A	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.4070	Added support for IPv6 and IPv4
	3.9.0300	Added support for unnumbered neighbors and Updated example

Example	switch (config router bgp 100)# neighbor 10.10.10.10 remove-private-as switch (config router bgp 200)# vrf default neighbor interface ethernet 1/1 remove-private-as
Related Commands	ip routing router bgp <as-number>
Notes	• This can only be used with external BGP (eBGP) peers • If the update has only private AS numbers in the AS path, BGP removes these numbers • If the AS path includes both private and public AS numbers, BGP does not remove the private AS numbers. This situation is considered a configuration error • If the AS path contains the AS number of the eBGP neighbor, BGP does not remove the private AS number • If the AS path contains confederations, BGP removes the private AS numbers only if they come after the confederation portion of the AS path

15.3.12.2.36 neighbor route-map

	neighbor <ip-address | peer-group-name | ethernet | port-channel> route-map <route-map-name> [in | out] no neighbor <ip-address | peer-group-name | ethernet | port-channel> route-map [route-map-name] [in | out] disable neighbor <ip-address | peer-group-name | ethernet | port-channel> route-map [route-map-name] [in | out] Configures route-map export or import to the peer either for a specific address family or for all (depending on the configuration context). The no form of the command removes map-route configuration and enables inheritance. The inheritance priority is as follows: a. Peer AFI-SAFI b. Peer c. Peer Group AFI-SAFI d. Peer Group The “disable” form of the command resets the route-map configuration to the default and disables inheritance.	
Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	route-map-name	Name of the route-map
	in | out	• in—sets route import to the peer for this AFI/SAFI • out—sets route export to the peer for this AFI/SAFI If no parameter is explicitly used, both in and out are configured.
	ethernet	interface ethernet <ifname>
	port-channel	interface port-channel <ifname>
Default	N/A	
Configuration Mode	config router bgp config router bgp address-family	
History	3.3.5006	
	3.3.5200	Updated notes and default
	3.4.1100	Added “out” parameter
	3.6.3004	Added note

	3.6.4070	Added support for IPv6 and IPv4
	3.6.8100	Added “config router bgp address-family” configuration mode
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	switch (config router bgp 100)# neighbor 10.10.10.10 route-map MyRouteMap in switch (config router bgp 65001 address-family l2vpn-evpn) # neighbor 192.168.3.2 route-map routeMapSample in switch (config router bgp 100 address-family ipv4-unicast) # neighbor 1.1.1.1 route-map sampleRoutemap in switch (config router bgp 200)# vrf default address-family ipv4-unicast neighbor interface ethernet 1/1 route-map r_map_test out	
Related Commands	neighbor <ip-address> remote-as <as-number> route-map <map-name> [deny | permit] [sequence-number] clear ip bgp {<ip-address> | all}	

Notes	• There are 3 possible ways of using the “disable” prefix:
	  • At the beginning of the command
	  switch (config) # disable router bgp 65001 address-family l2vpn-evpn neighbor 192.168.3.2 route-map
	  • After the “router bgp *”
	  switch (config) # router bgp 65001 disable address-family l2vpn-evpn neighbor 192.168.3.2 route-map
	  • After the “router bgp * address-family l2vpn-evpn”
	  switch (config) # router bgp 65001 address-family l2vpn-evpn disable neighbor 192.168.3.2 route-map
	• When inheritance is enabled (by default or when using the no form of the command), then if there is no peer AFI SAFI route-map configuration, then an assessment is made of whether a route-map was at the peer level or not. If yes, then it is taken. Otherwise, the OS continues looking to the peer group AFI SAFI, and then the peer group (if a peer is member of a peer group).
	  • Only one inbound route-map can be applied to a given neighbor
	  • If a new route-map is applied to a neighbor, it replaces the previous route map
	  • Changing a route-map only takes effect on routes received or sent after the change
	  • A route-map must already exist before a node is configured to use it	

15.3.12.2.37 neighbor no-route-map

	neighbor <ip-address> | <peer-group-name | ethernet | port-channel> no-route-map <route-map-name> [in|out] Unsets route-map for neighbor and prevents the system from inheriting this parameter’s group configuration.	
Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	route-map-name	Name of the route-map
	in | out	• in—sets route import to the peer for this AFI/SAFI • out—sets route export to the peer for this AFI/SAFI If no parameter is explicitly used, both in and out are configured.
	ethernet	interface ethernet <ifname>
	port-channel	interface port-channel <ifname>
Default	N/A	

Configuration Mode	config router bgp	
History	3.6.3004	
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 no-route-map` `switch (config router bgp 200)# vrf default address-family ipv4-unicast neighbor interface ethernet 1/1 no-route-map out`	
Related Commands	neighbor <ip-address> remote-as <as-number> route-map <map-name> [deny | permit] [sequence-number]	
Notes	BGP command "no-route-map" is deprecated and been replaced with the disable form of the BGP neighbor route-map command.	

15.3.12.2.38 neighbor route-reflector-client

	neighbor <ip-address | peer-group | ethernet | port-channel> route-reflector-client no neighbor <ip-address | peer-group | ethernet | port-channel> route-reflector-client disable neighbor <ip-address | peer-group | ethernet | port-channel> route-reflector-client Configures a given peer to be a reflector client of this router for this address-family. The no form of the command removes configuration and enables inheritance of AFI/SAFI route-reflector-client configuration from a peer group if this neighbor is member in one. The disable form of the command removes a given peer from being a reflector client of this router for this AFI/SAFI and disables configuration inheritance.	
Syntax Description	ip-address	Neighbor IP address
	peer-group	Peer group name
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	N/A	
Configuration Mode	config router bgp config router bgp address-family	
History	3.3.5006	
	3.3.5200	Updated notes and default
	3.6.3004	Added “disable” parameter
	3.6.4070	Added support for IPv6 and IPv4
	3.6.8100	Added “config router bgp address-family” configuration mode
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 route-reflector-client` `switch (config router bgp 200)# vrf default address-family ipv4-unicast neighbor interface ethernet 1/1 route-reflector-client`	
Related Commands		

Notes	There are 4 possible ways of using the “disable” prefix: • At the beginning of the command switch (config) # disable router bgp 65001 address-family l2vpn-evpn neighbor 192.168.3.2 route-reflector-client • At the end of the command switch (config) # router bgp 65001 address-family l2vpn-evpn neighbor 192.168.3.2 route-reflector-client disable • After the “router bgp *” switch (config) # router bgp 65001 disable address-family l2vpn-evpn neighbor 192.168.3.2 route-reflector-client • After the “router bgp * address-family l2vpn-evpn” switch (config) # router bgp 65001 address-family l2vpn-evpn disable neighbor 192.168.3.2 route-reflector-client

15.3.12.2.39 neighbor send-community

	neighbor <ip-address | peer group | ethernet | port-channel> send-community [extended] no neighbor <ip-address | peer group | ethernet | port-channel> send-community [extended] disable neighbor <ip-address | peer group | ethernet | port-channel> send-community [extended] Enables sending UPDATE messages to the peer containing BGP community attributes either for this address family or all relevant address-families. The no form of the command removes configuration and enables inheritance of send-community attribute configuration. The disable form of the command disables sending UPDATE messages containing BGP community attributes.	
Syntax Description	ip-address	Neighbor IP address
	peer_group	Peer group name
	extended	Enables sending UPDATE messages to the peer for this address family containing extended BGP community attributes
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	Enabled	
Configuration Mode	config router bgp config router bgp address-family	
History	3.4.0000	
	3.6.3004	Added “disable” parameter
	3.6.4070	Added support for IPv6 and IPv4
	3.6.8100	Added “config router bgp address-family” configuration mode
	3.9.0300	Added support for unnumbered neighbors and Updated example

Example	`switch (config router bgp 100)# neighbor 10.10.10.10 send-community`
	`switch (config router bgp 65001 address-family l2vpn-evpn) # neighbor 192.168.3.2 send-community`
	`switch (config router bgp 200)# vrf default address-family ipv4-unicast neighbor interface ethernet 1/1 send-community`	

Related Commands		

Notes	There are 4 possible ways of using the “disable” prefix: • At the beginning of the command switch (config) # disable router bgp 65001 address-family l2vpn-evpn neighbor 192.168.3.2 send-community • At the end of the command switch (config) # router bgp 65001 address-family l2vpn-evpn neighbor 192.168.3.2 send-community disable • After the “router bgp *” switch (config) # router bgp 65001 disable address-family l2vpn-evpn neighbor 192.168.3.2 send-community • After the “router bgp * address-family l2vpn-evpn” switch (config) # router bgp 65001 address-family l2vpn-evpn disable neighbor 192.168.3.2 send-community

15.3.12.2.40 neighbor shutdown

	neighbor <ip-address | peer-group-name | ethernet | port-channel> shutdown [disable] no neighbor <ip-address | peer-group-name | ethernet | port-channel> shutdown Disables BGP neighbor gracefully. The no form of the command enables BGP neighbor.	
Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	disable	Enables BGP neighbor and prevents the system from inheriting this parameter’s group configuration
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	Enabled	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.3.5200	Updated note
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 shutdown` `switch (config router bgp 200)# vrf default neighbor interface ethernet 1/1  shutdown`	
Related Commands		
Notes	• Disabling a neighbor terminates all its active sessions and removes associated routing information • A group’s shutdown immediately impacts every peer in this group, making them inherit this parameter	

15.3.12.2.41 neighbor soft-reconfiguration

	neighbor <ip-address | peer-group-name | ethernet | port-channel> soft-reconfiguration no neighbor <ip-address | peer-group-name | ethernet | port-channel> soft-reconfiguration Enables neighbor soft reconfiguration. The no form of the command disables neighbor soft reconfiguration.	
Syntax Description	peer-group-name	Peer group name
	ip-address	IP address of the neighbor
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	Enabled	
Configuration Mode	config router bgp	
History	3.6.4070	
Example	`switch (config router bgp 100)# neighbor 10.10.10.1 soft-reconfiguration`	
Related Commands		
Notes		

15.3.12.2.42 neighbor soft-reconfiguration inbound

	neighbor <ip-address | peer-group-name | ethernet | port-channel> soft-reconfiguration inbound no neighbor <ip-address | peer-group-name | ethernet | port-channel> soft-reconfiguration inbound Enables neighbor soft reconfiguration. The no form of the command disables neighbor soft reconfiguration.	
Syntax Description	ip-address	Neighbor IPv4 address
	peer-group-name	Peer group name
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	N/A	
Configuration Mode	config router bgp	
History	3.6.8100	
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	`switch (config router bgp 65001) # neighbor 192.168.3.2 soft-reconfiguration inbound` `switch (config router bgp 200)# vrf default neighbor interface ethernet 1/1 soft-reconfiguration inbound`	
Related Commands		
Notes	This command is mandatory to show received EVPN for this neighbor	

15.3.12.2.43 neighbor timers

	neighbor <ip-address | peer-group-name | ethernet | port-channel> timers <keep-alive> <hold-time> no neighbor <ip-address | peer-group-name | ethernet | port-channel> timers Configures the keepalive and hold times for a specified peer. The no form of the command resets the parameters to their default values.	
Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	keep-alive	The period between the transmission of consecutive keep-alive messages • Range: 1-3600 seconds • "0" means that keepalive is not sent and the connection does not expire • Explicitly configuring the default, "60", prevents the system from inheriting this parameter's group configuration
	hold-time	The period the switch waits for a keepalive or update message before it disables peering • Range: 3-7200 seconds • "0" means that keepalive is not sent and the connection does not expire • Explicitly configuring the default, "180", prevents the system from inheriting this parameter's group configuration
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	keep-alive—60 seconds hold-time—180 seconds	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.3.5200	Updated description
	3.6.3004	Updated "hold-time" and "keep-alive" parameter's syntax description
	3.6.4070	Added IPv6 and IPv4 support
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 timers 65 195` `switch (config router bgp 200)# vrf default neighbor interface ethernet 1/1  timers 10 20`	
Related Commands	neighbor <ip-address> remote-as <as-number>	
Notes	Hold time must be at least 3 seconds and should be three times longer than the keep-alive setting.	

15.3.12.2.44 neighbor transport connection-mode passive

	neighbor <ip-address | peer-group-name | ethernet | port-channel> transport connection-mode passive [disable] no neighbor <ip-address | peer-group-name | ethernet | port-channel> transport connection-mode passive Sets the TCP connection for the specified BGP neighbor or peer group to passive mode. The no form of the command sets the specified BGP neighbor or peer group to active connection mode.	
Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	disable	Sets the specified BGP neighbor or peer group to active connection mode and prevents the system from inheriting this parameter's group configuration
	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	TCP sessions initiated	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.3004	Added "disable" parameter
	3.6.4070	Added IPv6 and IPv4 support
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 transport connection-mode passive` `switch (config router bgp 200)# neighbor interface ethernet 1/1 transport connection-mode passive`	
Related Commands		
Notes	• When the peer's transport connection mode is set to passive, it accepts TCP connections for BGP, but does not initiate them • BGP peers in active mode can both accept and initiate TCP connections for BGP	

15.3.12.2.45 neighbor update-source

	neighbor <ip-address> update-source {ethernet <slot/port> | loopback <number> | port-channel <number> | vlan <vlan-id>} no neighbor <ip-address> update-source Configures the source-address for routing updates and to establish TCP connections with peers. The no form of the command disables configured source-address for routing updates and for TCP connection establishment with a peer.	
Syntax Description	ip-address	IP address of the neighbor
	ethernet <slot/port>	Ethernet interface
	loopback <number>	Loopback interface number

	vlan <vlan-id>	VLAN interface Range: 1-4094
	port-channel <number>	LAG interface Range: 1-4094
Default	BGP uses best local address	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.6.4070	Added IPv6 and IPv4 support
Example	`switch (config router bgp 100)# neighbor 10.10.10.2 update-source vlan 10`	
Related Commands		
Notes	If BGP update-source on neighbor is configured, the given interface's primary address is used as the source address. If BGP update-source configured on a peer group, the primary address is not guaranteed to be the source.	

15.3.12.2.46 neighbor no-update-source

	neighbor <ip-address> no-update-source Disables configured source-address for routing updates and for TCP connection establishment with a peer and prevents the system from inheriting this parameter's group configuration.
Syntax Description	N/A
Default	BGP uses best local address
Configuration Mode	config router bgp
History	3.6.3004
Example	`switch (config router bgp 100)# neighbor 10.10.10.2 no-update-source`
Related Commands	
Notes	

15.3.12.2.47 neighbor weight

	neighbor <ip-address | peer-group-name | ethernet | port-channel> weight <value> no neighbor <ip-address | peer-group-name | ethernet | port-channel> weight Assigns a weight attribute to paths from the specified neighbor. The no form of the command resets to default values.	
Syntax Description	ipv4_addr, ipv6_addr	IP address of the neighbor
	peer-group-name	Peer group name
	value	Weight value • Range: 0-65535 • Explicitly configuring a default value prevents the system from inheriting this parameter's group configuration

	ethernet	`interface ethernet <ifname>`
	port-channel	`interface port-channel <ifname>`
Default	Value is 32768 for router-originated paths and 0 for routes received through BGP	
Configuration Mode	config router bgp	
History	3.4.0000	
	3.6.4070	Added IPv6 and IPv4 support
	3.8.2000	Updated weight range
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	`switch (config router bgp 100)# neighbor 10.10.10.10 weight 100` `switch (config router bgp 200)# vrf default neighbor interface ethernet 1/1 weight 100`	
Related Commands		
Notes	• Weight values set through route map commands have precedence over neighbor weight command values • Other attributes are used only when all paths to the prefix have the same weight • A path's BGP weight is also configurable through route maps • When multiple paths to a destination prefix exist, the best-path selection algorithm prefers the path with the highest weight • Weight is the first parameter that the BGP best-path selection algorithm considers	

15.3.12.2.48 network

	network <ip_prefix length> no network <ip_prefix length> Configures a route for advertisement to BGP peers. The no form of the command removes the route from the BGP routes table, preventing its advertisement. The route is only advertised if the router has a gateway to the destination.	
Syntax Description	ip_prefix_length	A string that specific route map is assigned to the network.
	length	/24 or 255.255.255.0 format.
Default	N/A	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.3.5200	Updated description, syntax description, and notes
	3.6.4070	
	3.10.4300	Removed route-map-name option
Example	`switch (config router bgp 100)# network 10.10.10.0 /24 routemap`	
Related Commands		
Notes	• The parameters "ip-prefix" and "length" specify the route destination • The configuration zeros the host portion of the specified network address (e.g. 192.0.2.4/24 is stored as 192.0.2.0/24) • Address family is identified by the network address itself and not by the configuration command context	

15.3.12.2.49 redistribute

	[neighbor <peer_group>] redistribute {connected | static | ospf | ospf-internal | ospf-external} [<route-map>] no redistribute {connected | static | ospf} Enables redistribution of specified routes to the BGP domain. The no form of the command disables route redistribution from the specified source.	
Syntax Description	connected	Redistributes the direct routes
	static	Redistributes the user-defined (static) route
	peer_group	Route map name that modifies default route attributes
	ospf	Redistributes all routes learned by OSPF protocol
	ospf-internal	Redistributes all OSPF-learned routes which are marked as internal
	ospf-external	Redistributes all OSPF-learned routes which are marked as external
Default	No redistribution	
Configuration Mode	config router bgp	
History	3.2.1000	
	3.6.4070	
Example	`switch (config router bgp 100)# redistribute ospf`	
Related Commands		
Notes	• Multiple redistribution options can be applied • This command cannot be used with route-maps	

15.3.12.2.50 router-id

	router-id <ip-address> [force] no router-id [force] Configures a fixed router ID for BGP. The no form of the command removes the fixed router ID and restores the system default.	
Syntax Description	ip-address	IP Address identified the router ID
	force	Applies configuration while BGP is admin-up
Default	The Router ID is dynamically elected (no router-id). • If a loopback interface is configured, the router ID is set to the IP address of the loopback interface • If multiple loopback interfaces are configured, the router ID is set to the IP address of the loopback interface with the highest IP address • If no loopback interface is configured, the router ID is set to the highest IP address on a physical interface	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.6.3004	Added “force” parameter
Example	`switch (config router bgp 100)# router-id 10.10.10.10`	
Related Commands		

Notes	The IP address configured identifies the BGP speaker. The command triggers an automatic notification and session reset for the BGP neighbors.

15.3.12.2.51 route-map

	[neighbor <peer_group>] route-map <route_map_name> [in | out] no [neighbor <peer_group>] route-map <route_map_name> [in | out] Specifies a route map that will be applied in the given direction for specific address family. The no form of the command removes this configuration.	
Syntax Description	route_map_name	Name of a route map to apply
	in/out	Specifies in which direction the route map is applied. If nothing is given, route map is applied in both directions.
	peer_group	Peer group name
Default	N/A	
Configuration Mode	config router bgp	
History	3.6.4070	
Example	`switch (config router bgp 100)# route-map default in`	
Related Commands		
Notes		

15.3.12.2.52 timers bgp

	timers bgp <keep-alive> <hold> no timers bgp Configures the BGP keepalive and hold times. The no form of the command resets the parameters to their default settings.	
Syntax Description	keep-alive	Frequency with which keepalive messages are sent to its peer. Range: 1-3600 seconds. 0—no keep-alive messages are sent.
	hold	Interval after not receiving a keepalive message that a peer is declared dead. Range: 3-7200 seconds. 0—peer is held indefinitely regardless of keep-alive messages.
Default	Keepalive time—60 secs Hold time—180 secs	
Configuration Mode	config router bgp	
History	3.3.5006	
	3.3.5200	Updated syntax description, related commands and notes
	3.6.3004	This command is blocked
	3.9.2100	Updated example
Example	`switch (config router bgp 100)# vrf default neighbor 10.10.10.1 timers 3 10`	
Related Commands	ip routing neighbor timers router bgp <as-number> show ip bgp	

Notes	• Timer settings apply to every peer connection • The command “neighbor timers” configures the times on a specified peer connection • Hold time should be three times longer than the keepalive setting

15.3.12.2.53 vni

	vni <vni_value> no vni <vni_value> Create VNI on the router BGP. The no form of the command deletes VNI on the router BGP.	
Syntax Description	vni_value	Range: 1-16777214
Default	N/A	
Configuration Mode	config router bgp address-family l2vpn-evpn	
History	3.8.1000	
Example	`switch (config router bgp 100 vrf default address-family l2vpn-evpn) # vni 1000`	
Related Commands	router bgp <as-number>	
Notes	This command is irrelevant when using the enabled auto-create mode.	

15.3.12.2.54 vni rd

	vni <vni_value> rd <rd> no vni <vni_value> rd Configure route distinguisher to VNI. The no form of the command deletes route distinguisher configuration	
Syntax Description	vni_value	Range: 1-16777214
	rd	Route distinguisher address in the format "ip:value" Valid value: The valid IP and value needs to be between 0 to 65535
Default	N/A	
Configuration Mode	config router bgp address-family l2vpn-evpn	
History	3.8.1000	
Example	`switch (config router bgp 100 vrf default address-family l2vpn-evpn) # vni 1000 rd 2.3.4.5:15`	
Related Commands	vni	
Notes	This command is irrelevant when using the enabled auto-create mode.	

15.3.12.2.55 vni route-target

	vni <vni_value> route-target {both | import | export} <route_target> no vni <vni_value> route-target {both | import | export} Configure route target to VNI. The no form of the command deletes route distinguisher configuration.	
Syntax Description	vni_value	Range: 1-16777214

	route_target	Several route-targets can be configured for each VNI
		Valid ranges:
		• for ip: value should be [0..65535]
		• for as_num: values are:
		  • if as_num value is less or equal to 65535: value can be [0..4294967295]
		  • if as_num is more than 65535: value can be between 0 to 65535

Default	N/A	
Configuration Mode	config router bgp address-family l2vpn-evpn	
History	3.8.1000	
Example	`switch (config router bgp 100 vrf default address-family l2vpn-evpn) # vni 1000 route-target both 1.2.3.4:15`	
Related Commands	vni	
Notes	This command is irrelevant when using the enabled auto-create mode.	

15.3.12.2.56 vni auto-create

	vni auto-create no vni auto-create Enables auto-create mode on router bgp. The no form of the command disables auto-create mode on router bgp.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config router bgp address-family l2vpn-evpn	
History	3.8.1000	
	3.8.2200	Command was changed from "auto-create" to "vni auto-create"
Example	`switch (config router bgp 100 vrf default address-family l2vpn-evpn) # vni auto-create`	
Related Commands	vni	
Notes	Upon enabling auto-create, VNI is created automatically	

15.3.12.2.57 route-table prefix-list

	route-table prefix-list <prefix-list-name> <export | import> no route-table prefix-list <prefix-list-name> <export | import> Configure RTM policy for IPv4 or IPv6 address-family and bind it with a prefix-list in export direction from BGP RIB to routing table or import in the reverse direction. The no forms of the command removed the RTM policy for IPv4 or IPv6 address-family.	
Syntax Description	prefix-list-name	Specific prefix-list name
	export	Filtering from RIB to FIB
	import	Filtering from FIB to RIB

Default	N/A
Configuration Mode	config router bgp address-family
History	3.8.2100
Example	switch (config) # router bgp 1 address-family ipv4-unicast switch (config router bgp 1 address-family ipv4-unicast) # route-table prefix-list kuku import switch (config router bgp 1 address-family ipv4-unicast) # route-table prefix-list kuku export switch (config router bgp 1 address-family ipv4-unicast) # exit switch (config) # show ip bgp address-family ipv4-unicast Address family : IPv4 Maximum Path : 0/0 Redistribute : Total Neighbors : 1 Total peer-groups : 0 Total dynamic ranges : 0 Route table prefix list (import/export): list-name/list name
Related Commands	route-table prefix-list show ip bgp vrf address-family
Notes	Valid does both IPv4-unicast and IPv6-unicast

15.3.12.3 Show

15.3.12.3.1 show {ip | ipv6} bgp

Syntax Description
Default
Configuration Mode
History
Example

```
Output 1:
switch (config) # show ip bgp 192.168.100.0 /24

BGP table version: 22
Local router ID: 192.168.100.11

Status codes:
  s: suppressed
  d: damped
  h: history
  *: valid
  >: best
  i: internal
  r: RIB-failure
  S: Stale
  m: multipath
  b: backup-path
  x: best-external

Origin codes:
  i: IGP
  e: EGP
  ?: incomplete

------------------------------------------------------------------------
Network           Next Hop  Status  Metric   LocPrf     Weight      Path
------------------------------------------------------------------------
192.168.100.0/24  0.0.0.0   *>      0        100        32768       i


Output 2:
mtbc-baidu-01-2410 [standalone: master] (config) # show ip bgp BGP table version: 65 Local route
```

Related Commands
Notes

15.3.12.3.2 show ip bgp address-family

	show ip bgp address-family [vrf <vrf-name>] <l2vpn-evpn | ipv4-unicast | ipv6-unicast> [active] [detail] Displays address-family configuration.	
Syntax Description	l2vpn-evpn	Displays information about L2VPN-EVPN address family
	active	Displays active neighbors in that address family (configured, active or dynamic)
	detail	Displays detailed info about configuration and configured/active neighbors for the specified address-family
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4070	
	3.7.1000	Added “l2vpn-evpn” parameter and Updated example
	3.8.1000	Added output example for an updated address family configuration
	3.8.2100	Added RTM import/export policy
	3.8.2200	Updated output example for "show ip bgp address-family l2vpn-evpn"
Example		

Example output 1:

```
switch (config) # show ip bgp address-family l2vpn-evpn
Address family                          : L2VPN EVPN
Maximum Path                            : 0/0
Redistribute                            :
Total Neighbors                         : 0
Total peer-groups                       : 0
Total dynamic ranges                    : 0
Auto-Create VNI                         : Disable
Route table prefix list (import/export):
RD/RT Auto-Create                       : Disable


switch (config) # show ip bgp address-family l2vpn-evpn active
Address family               : L2VPN EVPN
Networks                     :
maximum-path                 : 0/0
redistribute                 : -
Total neighbors              : 2
Total peer-groups            : 0
Total dynamic ranges         : 0
```

```
switch (config) # show ip bgp address-family l2vpn-evpn detail

Address family    : L2VPN EVPN
Maximum Path      : 0/0
Redistribute      :
Total Neighbors   : 1

Neighbors:
  ------------------------------------------------------------------------------------------
  Neighbor   V   AS      MsgRcvd   MsgSent   TblVer   InQ   OutQ   Up/Down   State/PfxRcd
  ------------------------------------------------------------------------------------------
  1.1.1.1    4   65002   0         1         6        0     0      Never     ACTIVE/0

Total peer-groups         : 1
Peer Group                : peer
Total dynamic ranges      : 0
Auto-Create VNI           : Disable
-------------------------------------------------------------------------------
VNI          Vlan         Route Distinguisher   Route Target
-------------------------------------------------------------------------------
1000         5            1.2.3.4:3             None
```

```
Example output 2:
switch (config) # show ip bgp address-family ipv4-unicast detail
Address family : IPv4
Maximum Path : 0/0
Redistribute :
Total Neighbors: 1

Neighbors:
  ------------------------------------------------------------------------------------------
  Neighbor   V   AS     MsgRcvd    MsgSent   TblVer   InQ   OutQ   Up/Down   State/PfxRcd
  ------------------------------------------------------------------------------------------
  3.3.3.3    4   200    0          0         1        0     0      Never     IDLE/0

Total peer-groups : 1
Peer Group : basim_ipv4
Total dynamic ranges: 0
Address family configuration:
Next hop unchanged: Enable


Example output 3:
```

```
switch (config) # show ip bgp address-family ipv4-unicast

Address family                          : IPv4
Maximum Path                            : 0/0
Redistribute                            :
Total Neighbors                         : 1
Total peer-groups                       : 0
Total dynamic ranges                    : 0
Route table prefix list (import/export): a-list/a-list
```

Related Commands	

Notes	

15.3.12.3.3 show ip bgp community

	show ip bgp [vrf <vrf-name>] community <comm1> <comm2> ... <commn> [exact] [detail] Displays information about the BGP routes (RIB) filtered according to communities.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.0000	
	3.6.4070	Added support for IPv6

Example

```
switch (config) # show ip bgp community 100:1

BGP table version is 8, local router ID is 3.5.7.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
              r RIB-failure, S Stale, m multipath, b backup-path, x best-external
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network            Next Hop            Metric     LocPrf     Weight Path
*>  3.4.3.11/32        0.0.0.0                  0          0      32768 i
*>  3.5.7.88/32        0.0.0.0                  0          0      32768 i
*>  3.5.7.99/32        0.0.0.0                  0          0      32768 i

switch (config) # show ip bgp community 100:1 exact
BGP table version is 8, local router ID is 3.5.7.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
              r RIB-failure, S Stale, m multipath, b backup-path, x best-external
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network            Next Hop            Metric     LocPrf     Weight Path
*>  3.4.3.11/32        0.0.0.0                  0          0      32768 i
*>  3.5.7.99/32        0.0.0.0                  0          0      32768 i

switch (config) # show ip bgp community 100:1
BGP table version is 8/20, local router ID is 3.5.7.4
Status codes: * valid, > best, i - internal, m multipath
Origin codes: i - IGP, e - EGP, ? - incomplete
Network Next Hop Metric LocPrf Weight Path
*> 2001::0/64 2001:1::1 0 0 32768 i
```

Related Commands	show ip bgp
Notes	

15.3.12.3.4 show ip bgp evpn

	show ip bgp [vrf <vrf-name>] [neighbors <ip | peer-group | ethernet | port-channel> [received | advertised]] evpn [route-type <type> | community {<aa:nn> | <number>} | extcommunity route-target {<aa:id> | <aa.bb:id> | <ip:id>} | extcommunity router-mac <mac-address> | vni <value> | rd <rd>] [detail] Displays BGP EVPN routes received from all neighbors in specified VRF or the VRF currently under context.	
Syntax Description	ipv4_addr	Neighbor IP address
	peer_group	Peer group name
	ethernet	interface ethernet <ifname>

	port-channel	interface port-channel <ifname>
	route-type	Possible values: • auto-discovery—Ethernet Auto-discovery Route • mac-ip—MAC/IP Advertisement Route • imet—Inclusive Multicast Ethernet Tag Route • ethernet-segment—Ethernet Segment Route • ip-prefix—IP Prefix Route
	community	<aa:nn>—community number <number>—community number
	extcommunity route-target	Filters by route target • <aa:id>—Route Target (asplain) • <aa.bb:id>—Route Target (asdot) • <ip:id>—Rout Target (IP)
	extcommunity router-mac	Filters by router MAC
	vni	VNI value Range: 1-16777215
	rd	Filters by route target • <aa:id>—Route Target (asplain) • <aa.bb:id>—Route Target (asdot) • <ip:id>—Rout Target (IP)
	detail	Shows additional information about BGP route
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8100	
	3.8.2200	• Added "show ip bgp evpn detail" output • Replaced auto-completion of “route-type” with string keywords instead on numbers
	3.9.0300	Adding vni attribute to "show ip bgp evpn detail" for imet routes and added example
	3.9.1000	Added ability to select several attributes for filtering output routes
Example		

```
switch (config) # show ip bgp evpn

------------------------------------------------------------------------------------------------------------------
----------------------------
RD                    Type              Data                              Next Hop         Metric
LocPrf      Weight      Path
------------------------------------------------------------------------------------------------------------------
----------------------------
2.3.4.5:5             mac-ip            00:bb:cc:dd:ee:ff                 2.3.4.5          0           100
0           ?
2.3.4.5:6             mac-ip            00:aa:bb:cc:dd:ee                 2.3.4.5          0           100
0           ?
1.2.3.4:5             imet              1.2.3.4                           1.2.3.4          0           100
0           ?
1.2.3.4:6             imet              1.2.3.4                           1.2.3.4          0           100
0           ?
2.3.4.5:5             imet              2.3.4.5                           2.3.4.5          0           100
0           ?
2.3.4.5:6             imet              2.3.4.5                           2.3.4.5          0           100
0           ?



switch (config) # show ip bgp evpn vni 1000

------------------------------------------------------------------------------------------------------------------
----------------------------
RD                    Type              Data                              Next Hop         Metric
LocPrf      Weight      Path
------------------------------------------------------------------------------------------------------------------
----------------------------
2.3.4.5:5             mac-ip            00:bb:cc:dd:ee:ff                 2.3.4.5          0           100
0           ?
1.2.3.4:5             imet              1.2.3.4                           1.2.3.4          0           100
0           ?
2.3.4.5:5             imet              2.3.4.5                           2.3.4.5          0           100
0           ?

switch (config) # show ip bgp evpn vni 1000 route-type mac-ip

------------------------------------------------------------------------------------------------------------------
----------------------------
RD                    Type              Data                              Next Hop         Metric
LocPrf      Weight      Path
------------------------------------------------------------------------------------------------------------------
----------------------------
2.3.4.5:5             mac-ip            00:bb:cc:dd:ee:ff                 2.3.4.5          0           100
0           ?

switch (config) # show ip bgp evpn vni 1000 route-type mac-ip detail

1 paths for mac-ip 00:bb:cc:dd:ee:ff   Route Distinguisher: 2.3.4.5:5:
  route:
    next hop          : 2.3.4.5
    neighbor ip       : 1.1.1.2
    router id         : 2.3.4.5
    metric            : 0
    weight            : 0
    local pref        : 100
    origin            : incomplete
    Extended Community: 100:268436456(Route-Target-AS)
    Extended Community: tunnelTypeVxlan(TunnelEncap)
    flags             : valid, best
    esi               : 00:00:00:00:00:00:00:00:00:00
    vni               : 1000
    path              :
    ethernet tag id   :
```

Related Commands	
Notes	

15.3.12.3.5 show ip bgp evpn summary

	show ip bgp [vrf <vrf>] evpn summary Displays some basic statistics about BGP per VRF only for neighbors who support L2EVPN AF.	
Syntax Description	vrf	Name of VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8100	
Example		

```
switch (config) # show ip bgp evpn summary

VRF name                                : vrf-default
BGP router identifier                   : 192.168.5.1
local AS number                         : 65001
BGP table version                       : 2
Main routing table version              : 2
IPV4 Prefixes                           : 0
IPV6 Prefixes                           : 0
L2VPN EVPN Prefixes                     : 1

------------------------------------------------------------------------------------------
Neighbor      V   AS      MsgRcvd   MsgSent   TblVer   InQ   OutQ   Up/Down      State/PfxRcd
------------------------------------------------------------------------------------------
192.168.3.2   4   65002   25        29        2        0     0      0:00:11:10   ESTABLISHED/1
192.168.5.2   4   65003   24        28        2        0     0      0:00:11:17   ESTABLISHED/0
```

Related Commands	
Notes	

15.3.12.3.6 show ip bgp exceptions

	show ip bgp exceptions Displays all the bgp exceptions.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.9.0300

Example	`switch (config)# show ip bgp exceptions`
	`-----------------------------------------------------------------------`
	`Type  Origin  Description`
	`-----------------------------------------------------------------------`
	`neighbor  ethernet 1/11  Interface doesn't exist`

Related Commands	router bgp neighbor interface
Notes	

15.3.12.3.7 show ip bgp neighbors

	show ip bgp [vrf <vrf-name>] neighbors <ip-address | ethernet | port-channel> Displays summaries information about all BGP neighbors.

Syntax Description	vrf	VRF name
	ip	Neighbor IPv4 address
	ethernet	interface ethernet <ifname>
	port-channel	interface port-channel <ifname>
	ifname	Interface number (Ethernet or port-channel number)
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.7.1000	Updated example
	3.8.2200	Updated xample to reflect the new "Enhanced Route Refresh" display
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example		

Output 1:

```
switch (config) # show ip bgp neighbors 192.168.10.2
  BGP neighbor: 192.168.10.2, remote AS: 100, link: internal:
  BGP version : 4
  Configured hold time in seconds : 180
  keepalive interval in seconds (configured) : 60
  keepalive interval in seconds (established with peer): 60
  Minimum holdtime from neighbor in seconds : 180
  Peer group :
Neighbor configuration:
    ------------------------------------------------------------------------
    Configuration                   IPV4 Unicast  IPV6 Unicast  L2VPN EVPN
    ------------------------------------------------------------------------
    Configured AFI SAFI             Enabled       Disabled      Enabled
    Send Community                  Enabled       Enabled       Enabled
    Send Extended Community         Enabled       Enabled       Enabled
    Route Reflection                Disabled      Disabled      Disabled
    Next Hop Unchanged              Disabled      Disabled      Enabled
    Extended next hop IPv4          Disabled      Disabled      Disabled
Neighbor capabilities:
    Route Refresh               : advertise and received
    Enhanced Route Refresh      : advertise and received
    Soft Reconfiguration        : Disabled
    Graceful Restart Capability: advertise and received
    Address family IPv4 Unicast: advertise and received
    Address family IPv6 Unicast: n/a
    Address family L2VPN EVPN   : advertise and received
    Extended next hop IPv4      : n/a
Message statistics:
    InQ depth : 0
    OutQ depth: 0
    ------------------------------------------------
    Parameter             Sent         Rcvd
    ------------------------------------------------
    Opens                 1            1
    Notification          0            0
    Updates               4            4
    Keepalives            9            9
    Refreshes             0            0
    Total                 14           14
Default minimum time between advertisement runs in seconds: 30
IPV4 Unicast:
    ------------------------------------------------
    Prefix activity       Sent         Rcvd
    ------------------------------------------------
    Prefixes Current      1            1
    Prefixes Total        1            1
    Implicit Withdraw     0            0
    Explicit Withdraw     0            0
    Used as bestpath      n/a          1
    Used as multipath     n/a          n/a
    ------------------------------------------------------------
    Local Policy Denied Prefixes    Outbound     Inbound
    ------------------------------------------------------------
    Total                           2            0
```

```
L2VPN EVPN:
    ---------------------------------------------
    Prefix activity       Sent         Rcvd
    ---------------------------------------------
    Prefixes Current      1            1
    Prefixes Total        1            1
    Implicit Withdraw     0            0
    Explicit Withdraw     0            0
    Used as bestpath      n/a          1
    Used as multipath     n/a          n/a

    -----------------------------------------------------------
    Local Policy Denied Prefixes    Outbound     Inbound
    -----------------------------------------------------------
    Total                           2            0
Connection Information:
 Connections established : 1
 Dropped : 0
 Last Reset : 0:00:06:59
 Last Drop Reason : 6 (2)
 Maximum hops to external BGP neighbor: 255
 Connection State : ESTABLISHED
 Local host : 192.168.1.1
 Local port : 56794
 Foreign host : 192.168.10.2
 Remote port : 179


Output 2:

switch (config) # show ip bgp neighbors ethernet 1/1

BGP neighbor: ethernet 1/1, remote AS: 65002, link: external:
BGP version                                 : 4
  Configured hold time in seconds           : 180
  keepalive interval in seconds             : 60
  Minimum holdtime from neighbor in seconds: 90
  Peer group                                :

  Neighbor configuration:
    ------------------------------------------------------------------------
    Configuration                   IPV4 Unicast  IPV6 Unicast  L2VPN EVPN
    ------------------------------------------------------------------------
    Configured AFI SAFI             Enabled       Disabled      Enabled
    Send Community                  Disabled      Disabled      Disabled
    Send Extended Community         Disabled      Disabled      Disabled
    Route Reflection                Disabled      Disabled      Disabled
    Next Hop Unchanged              Disabled      Disabled      Disabled

  Neighbor capabilities:
    Route Refresh              : advertise and received
    Enhanced Route Refresh     : advertise and received
```

```
    Soft Reconfiguration       : Disabled
    Graceful Restart Capability: advertise
    Address family IPv4 Unicast: advertise and received
    Address family IPv6 Unicast: n/a
    Address family L2VPN EVPN  : advertise and received
  Message statistics:
    InQ depth : 0
    OutQ depth: 0

    ---------------------------------------
    Parameter             Sent         Rcvd
    ---------------------------------------
    Opens                 1            1
    Notification          0            0
    Updates               3            2
    Keepalives            12           11
    Refreshes             0            0
    Total                 16           14

    Default minimum time between advertisement runs in seconds: 30

  L2VPN EVPN:
    ---------------------------------------------
    Prefix activity       Sent         Rcvd
    ---------------------------------------------
    Prefixes Current      2            2
    Prefixes Total        2            2
    Implicit Withdraw     0            0
    Explicit Withdraw     0            0
    Used as bestpath      n/a          2
    Used as multipath     n/a          n/a

    -----------------------------------------------------
    Local Policy Denied Prefixes    Outbound     Inbound
    -----------------------------------------------------
    Total                           0            0

  Connection Information:
    Connections established              : 4
    Dropped                              : 1
    Last Reset                           : 0:00:03:22
    Last Drop Reason                     : 6 (2)
    Maximum hops to external BGP neighbor: 255
    Connection State                     : ESTABLISHED
    Local host                           : 192.168.2.1
    Local port                           : 179
    Foreign host                         : 192.168.2.2
    Local Port                           : 50394


Output 3:
switch (config) # show ip bgp neighbors

BGP neighbor: 192.168.2.2, remote AS: 65001, link: internal:
  BGP version                               : 4
  Configured hold time in seconds           : 180
  keepalive interval in seconds             : 60
  Minimum holdtime from neighbor in seconds: 90
```

Related Commands	
Notes	

15.3.12.3.8 show ip bgp neighbors advertised/received address-family

	show ip bgp neighbors <neigh_ip | ethernet | port-channel> <advertised | received> <ipv4-unicast | ipv6-unicast> Displays advertised/received BGP routes for a specific address-family per neighbor.	
Syntax Description	neigh_ip	Neighbor IP address
	ethernet	interface ethernet <ifname>

	port-channel	interface port-channel <ifname>
Default	N/A	
Configuration Mode	Any command mode	
History	3.8.2200	
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example	Output 1: switch (config) # show ip bgp neighbors 192.168.7.2 advertised ipv4-unicast BGP table version: 2 Local router ID : 192.168.1.1 Status codes: s: suppressed d: damped h: history *: valid >: best i: internal r: RIB-failure S: Stale m: multipath b: backup-path x: best-external Origin codes: i: IGP e: EGP ?: incomplete --- Network Next Hop Status Metric LocPrf Weight Path --- 192.168.1.1/32 192.168.7.1 i* 0 100 32768 i Output 2: switch (config) # show ip bgp neighbors interface ethernet 1/17 advertised ipv4-unicast BGP table version: 65 Local router ID : 22.1.1.1 ... --- Network Next Hop Status Metric LocPrf Weight Path --- 17.1.1.0/24 Eth1/17 * 0 0 32768 300 ? ...	
Related Commands		
Notes	• In order to use received option, user must first configure soft-reconfiguration-inboud as follows: switch (config) # router bgp 100 neighbor 192.168.7.2 soft-reconfiguration inbound • Received option "shows BGP routes" shows all received routes before applying policies • Advertised option shows BGP routes after applying policies.	

15.3.12.3.9 show ip bgp neighbors received

	show ip bgp neighbors <ip-address | ethernet | port-channel> received [<ip-address> [<mask>] [longer-prefixes]] Displays BGP summary information.

Syntax Description	ip-address	Neighbor IP address
	mask	Mask length
	longer-prefixes	Displays the routes to the specified destination and any routes to a more specific destination (only available if both IP and mask are specified)
	ethernet	interface ethernet <ifname>
	port-channel	interface port-channel <ifname>
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.7.1000	Updated example
	3.8.1000	Updated example
	3.9.0300	Added support for unnumbered neighbors and Updated example

Example

```
Output 1:
switch (config) # show ip bgp neighbors 192.168.3.2 received

BGP table version: 16
local router ID : 192.168.1.1

Status codes:
 s: suppressed
 d: damped
 h: history
 *: valid
 >: best
 i: internal
 r: RIB-failure
 S: Stale
 m: multipath
 b: backup-path
 x: best-external

Origin codes:
 i: IGP
 e: EGP
 ?: incomplete

------------------------------------------------------------------
Network        Next Hop     Status    Metric  LocPrf  Weight    Path
------------------------------------------------------------------
94.0.0.0/24    192.168.3.2   *>       0       100     0       100 i


Output 2:
switch (config) # show ip bgp neighbors interface ethernet 1/17 received

BGP table version: 65
local router ID  : 22.1.1.1
...
------------------------------------------------------------------------
Network        Next Hop    Status    Metric   LocPrf   Weight    Path
------------------------------------------------------------------------
17.1.1.0/24    Eth1/17     0         100      0        23 ?
...
```

Related Commands	
Notes	

15.3.12.3.10 show ip bgp neighbors received detail

	show ip bgp neighbors <ip-address> [received] [<ip-address> [<mask> [longer-prefixes]]] detail Displays detailed information on routes received from neighbors.	
Syntax Description	ip-address	Neighbor IP address. Provide optionally to display routes received from specified neighbor.
	mask	Mask length. Displays routes received from specified neighbor filtered by the specified network.
	longer-prefixes	Displays routes received from specified neighbor filtered by the specified prefix and longer
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.7.1000	Updated example
Example		

```
switch (config)# show ip bgp 192.168.100.0 /24 longer-prefixes detail

BGP routing table entry for: 192.168.100.0/24
Version                    : 22
Paths                      : (1, best: #1)

Local Connected:
  Origin    : IGP
  metric    : 0
  localpref : 100
  weight    : 32768
  Attributes: valid, best
switch (config)# show ip bgp 192.168.100.0 /24 detail

BGP routing table entry for: 192.168.100.0/24
Version                    : 22
Paths                      : (1, best: #1)

Local connected:
  0.0.0.0 from 0.0.0.0 (192.168.100.11):
    Origin    : IGP
    metric    : 0
    localpref : 100
    weight    : 32768
    Attributes: valid, sourced, best
```

Related Commands	
Notes	

15.3.12.3.11 show ip bgp paths

	show ip bgp paths [vrf <vrf-name>] [ipv4 | ipv6] Displays summary of all AS paths and for prefixes for specific address family.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.6.4070	Added support for IPv4 and IPv6

	3.9.2300	Updated example
Example	switch (config) # show ip bgp paths Path Metric Refcount 4 50 100 0 1 2 50 100 0 1 4 40 0 1 12 50 100 0 1 2 0 1 2 20 0 1	
Related Commands		
Notes		

15.3.12.3.12 show ip bgp peer-group

	show ip bgp peer-group [vrf <vrf-name>] [peer-group-name] <ipv4-unicast |ipv6-unicast> Displays information about peer groups and configuration, filtered per address family.	
Syntax Description	peer-group-name	Displays information about a specific peer-group.
Default	N/A	
Configuration Mode	Any command mode	
History	3.4.0000	
	3.6.8100	Updated example
	3.7.1000	Updated example
Example		

```
switch (config) # show ip bgp peer-group peerGrp1

Name                     : peerGrp1
Hold time                : 180
Keep-alive               : 60
Max prefix               : 100000
Weight                   : 0
Export local preferences: 100
Import local preferences: 100
Status Down              : no
EBGP Multihop            : 1
Next Hop Self            : no
Soft Reconfiguration     : no
Next Hop Peer            : no
Remove Private AS        : no
Transport Mode           : no
Password                 : no
Local AS                 : 0
No Prepend               : no
Replace AS               : no
Soft Reconfiguration     : Disabled

-------------------------------------------------------------------------
Configuration                   IPV4 Unicast  IPV6 Unicast  L2VPN EVPN
-------------------------------------------------------------------------
Configured AFI SAFI             Disabled      Disabled      Disabled
Send Community                  Disabled      Disabled      Disabled
Send Extended Community         Disabled      Disabled      Disabled
Route Reflection                Disabled      Disabled      Disabled
Next Hop Unchanged              Disabled      Disabled      Disabled

-------------------------------------------------------------------------------------------------------
Neighbor          V     AS     MsgRcvd   MsgSent   TblVer    InQ    OutQ   Up/Down       State/PfxRcd

-------------------------------------------------------------------------------------------------------
192.168.2.2       4     65001 355        413       7         0      0      0:00:00:26    ESTABLISHED/2
```

Related Commands	
Notes	

15.3.12.3.13 show ip bgp summary

	show ipv6 bgp {<id> | all} [vrf <vrf-name>] summary Displays BGP summary for IPv6 addresses.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.6.4070	Added support for IPv6
	3.9.0300	Updated example to reflect support of BGP unnumbered feature
Example		

```
Output 1:
switch (config) # show ip bgp summary

BGP router identifier 3.5.7.4, local AS number 4
BGP table version is 70/120, main routing table version 70/96
BGP using 26308 total bytes of memory
BGP activity  37/8  IPv4 prefixes,  37/8  IPv6 prefixes, 37/4  paths

-------------------------------------------------------------------------------------------------
Neighbor  V   AS   MsgRcvd  MsgSent InQ  OutQ Up/Down     State/PfxRcd
-------------------------------------------------------------------------------------------------
2001::1   4   7    3        9       0    0     0:00:00:48 ESTABLISHED/total number of prefixes


Output 2:
switch (config) # show ip bgp vrf default summary

VRF name                 : default
BGP router identifier    : 22.1.1.1
local AS number          : 300
BGP table version        : 31
Main routing table version: 31
IPV4 Prefixes            : 8
IPV6 Prefixes            : 2
L2VPN EVPN Prefixes      : 0

-------------------------------------------------------------------------------------------
Neighbor    V   AS   MsgRcvd   MsgSent  TblVer  InQ   OutQ  Up/Down       State/PfxRcd
-------------------------------------------------------------------------------------------
Eth1/17     4   23   378       377      31      0     0     0:05:05:14    ESTABLISHED/6
17.1.1.23   4   23   79        80       31      0     0     0:01:04:34    ESTABLISHED/4
2323::1     4   100  0         0        31      0     0     Never         IDLE/0
```

Related Commands	
Notes	

15.3.12.3.14 show ip bgp update-group

	show ip bgp update-group <neighbor ip address | ethernet | port-channel> Displays update-group information for all neighbors.	
Syntax Description	ethernet	interface ethernet <ifname>
	port-channel	interface port-channel <ifname>
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4070	
	3.7.1000	Updated example
	3.9.0300	Added support for unnumbered neighbors and Updated example
Example		

```
switch (config)# show ip bgp update-group 192.168.2.2

Update-group for neighbor: 192.168.2.2
BGP router identifier    : 192.168.2.1
local AS number          : 65001
BGP table version        : 7


------------------------------------------------------------------------------------------------
Neighbor          V    AS       MsgRcvd   MsgSent   TblVer    InQ    OutQ   Up/Down       State/PfxRcd

------------------------------------------------------------------------------------------------
192.168.2.2     4   65001    368       428       7         0      0     0:00:06:30   ESTABLISHED/2


r-mgtswd-270 [standalone: master] (config) # show ip bgp update-group

Update-group                                      : 5
BGP version                                       : 4
Address Family                                    : IPv4 Unicast
Minimum time between advertisements runs in seconds: 30

Has 1 members:
  192.168.2.2

Update-group                                      : 6
BGP version                                       : 4
Address Family                                    : L2VPN EVPN
Minimum time between advertisements runs in seconds: 30

Has 1 members:
  192.168.2.2


switch (config) #   show  ip bgp update-group interface ethernet 1/1
Update-group for neighbor: ethernet 1/1
BGP router identifier    : 2.2.2.2
local AS number          : 200
BGP table version        : 1


------------------------------------------------------------------------------------------------------------
--------
Neighbor          V    AS           MsgRcvd   MsgSent   TblVer    InQ    OutQ   Up/Down       State/PfxRcd

------------------------------------------------------------------------------------------------------------
--------
Eth1/1            4    100          6         7         1         0      0      0:00:03:23    ESTABLISHED/
0
```

Related Commands	
Notes	

15.3.12.3.15 show ip bgp vrf summary

	show ip bgp vrf {<vrf-name> | all} summary Displays BGP summary info for all or specified VRFs.	
Syntax Description	vrf-name	Displays BGP summary for specified VRF
	all	Displays BGP summary for all VRFs
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6000	
	3.6.8100	Updated example
Example		

```
switch (config)# show ip bgp summary

VRF name                                : vrf-default
BGP router identifier                   : 1.1.1.2
local AS number                         : 65001
BGP table version                       : 3
Main routing table version              : 3
IPV4 Prefixes                           : 0
IPV6 Prefixes                           : 0
L2VPN EVPN Prefixes                     : 2

-----------------------------------------------------------------------------------------
Neighbor   V    AS      MsgRcvd   MsgSent   TblVer   InQ   OutQ   Up/Down      State/PfxRcd
-----------------------------------------------------------------------------------------
1.1.1.1    4    65002   25        29        3        0     0      0:00:10:38   ESTABLISHED/2
1.1.1.5    4    100     0         0         3        0     0      Never        IDLE/0
```

Related Commands	
Notes	

15.3.12.4 IP AS-Path Access-List

15.3.12.4.1 ip as-path access-list

	ip as-path access-list <list-name> {permit | deny} <reg-exp> [any | egp | igp | incomplete] no ip as-path access-list <list-name> Creates an access list to filter BGP route updates. The no ip as-path access-list command deletes the named access list.	
Syntax Description	list-name	The name for the access list
	permit	Permits access for a matching condition
	deny	Denies access for a matching condition
	reg-exp	POSIX-compliant regular expression that is used to specify a pattern to match against an input string **Note:** The AS path is a comma-separated list of decimal AS numbers
	any	Any route type
	egp	External BGP routes
	igp	Internal BGP routes
	incomplete	Routes marked as “Incomplete”
Default	N/A	
Configuration Mode	config	
History	3.4.0000	
	3.9.2400	Updated "reg-exp" syntax description
Example	`switch (config)# ip as-path access-list mylist permit`	
Related Commands		
Notes	If access list_name does not exist, this command creates it. If it already exists, this command appends statements to the list.	

15.3.12.4.2 show ip as-path access-list

	show ip as-path access-list [list-name] Presents defined as-path access lists	
Syntax Description	list-name	Displays a specific prefix-list
Default	N/A	
Configuration Mode	config	
History	3.4.0000	
Example	`switch (config)# show ip as-path access-list mylist`	
Related Commands		
Notes		

15.3.12.5 IP Community-List

15.3.12.5.1 ip community-list standard

	ip community-list standard <list-name> {deny | permit} <list-of-communities> no ip community-list standard <list-name> Adds a standard entry to a community-list. The no form of the command deletes the specified community list.	
Syntax Description	list-name	The name for the community list
	permit	Permits access for a matching condition
	deny	Denies access for a matching condition
	list-of-communities	List of standard communities: • <aa:nn> • <number> • internet • local-AS • no-advertise • no-export
Default	N/A	
Configuration Mode	config	
History	3.4.0000	
Example	`switch (config)# ip community-list standard mycommunity permit 1:2 3:4`	
Related Commands		
Notes	A BGP community access list filters route maps that are configured as BGP communities. The command uses regular expressions to name the communities specified by the list.	

15.3.12.5.2 ip community-list expanded

	ip community-list expanded <list-name> {deny | permit} <reg-exp> no ip community-list expanded <list-name> Adds a regular expression entry to a community-list. The no form of the command deletes the specified community list.	
Syntax Description	list-name	Configures a named standard community list
	permit	Permits access for a matching condition
	deny	Denies access for a matching condition
	reg-exp	Regular expression that is used to specify a pattern to match against an input string
Default	N/A	
Configuration Mode	config	
History	3.4.0000	
Example	`switch (config)# ip community-list expanded mycommunity permit 1:[0-9]+`	
Related Commands		
Notes	A BGP community access list filters route maps that are configured as BGP communities. The command uses regular expressions to name the communities specified by the list.	

15.3.12.5.3 show ip community-list

	show ip community-list [community-list-name] Displays the defined community lists.	
Syntax Description	community-list-name	An optional parameter to display only the specified list
Default	N/A	
Configuration Mode	config	
History	3.4.0000	
Example	`switch (config)# show ip community-list mycommunity`	
Related Commands		
Notes	A BGP community access list filters route maps that are configured as BGP communities. The command uses regular expressions to name the communities specified by the list.	

15.3.13 BGP Monitoring Protocol

BGP monitoring protocol (BMP) is defined in RFC 7854, and is used to monitor BGP sessions. BMP is used to exchange BGP speaker status with a BMP collector. Usually, this speaker installs a number of BGP sessions with peers and one (or more) BMP sessions with a collector. The BGP speaker updates

the BMP server with the data received from its protocol, concerning changes in its peer sessions, and periodically sends out BGP statistics.

15.3.13.1 BMP Commands

15.3.13.1.1 protocol bmp

	protocol bmp no protocol bmp Enables BMP. The no form of the command disables BMP.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.7.1100
Example	`switch (config)# protocol bmp`
Related Commands	
Notes	• BMP commands are not executed when protocol BMP is disabled • Running protocol BMP when “no ip l3” is configured is not possible

15.3.13.1.2 ip bmp server

	ip bmp [vrf <vrf name>] server <id> no ip bmp [vrf <vrf name>] server <id> Creates a BMP server, up to three servers per VRF. The no form of the command removes BMP server configuration.	
Syntax Description	id	BMP server id: 1-3
	vrf name	The default is “default VRF”
Default	N/A	
Configuration Mode	config	
History	3.7.1100	
Example	`switch (config)# ip bmp server 1`	
Related Commands		
Notes		

15.3.13.1.3 ip bmp server activate

	ip bmp [vrf <vrf name>] server <id> activate no ip bmp [vrf <vrf name>] server <id> activate Activates BMP server. The no form of the command deactivates the BMP server.
Syntax Description	N/A
Default	N/A
Configuration Mode	config

History	3.7.1100
Example	switch (config)# ip bmp server 1 activate switch (config)# ip bmp server 1 vrf default activate
Related Commands	
Notes	

15.3.13.1.4 ip bmp server stats-reporting-period

	ip bmp [vrf <vrf name>] server <id> stats-reporting-period <seconds> no ip bmp [vrf <vrf name>] server <id> stats-reporting-period <seconds> Configures statistics reporting period. The no form of the command removes statistics reporting period configuration.	
Syntax Description	Seconds	Reporting period Range: 1-600 Default: 30
Default	N/A	
Configuration Mode	config	
History	3.7.1100	
Example	switch (config)# ip bmp server 1 stats-reporting-period 111	
Related Commands		
Notes	It is not possible to update a server's stats-reporting-period while the server is active	

15.3.13.1.5 ip bmp server address port

	ip bmp [vrf <vrf name>] server <id> address <address> port <port> no ip bmp [vrf <vrf name>] server <id> address <address> port <port> Configures an address for BMP server. The no form of the command removes address for BMP server.	
Syntax Description	address	IPv4 or IPv6 server address
	port	TCP port to connect
Default	N/A	
Configuration Mode	config	
History	3.7.1100	
Example	switch (config)# ip bmp server 1 address 1.1.1.1 port 11 switch (config)# ip bmp server 1 vrf vrf-default address 7.7.7.7 port 5000	
Related Commands		
Notes	It is not possible to update a server's address while the server is active	

15.3.13.1.6 show ip bmp

	show ip bmp [vrf <vrf name>] [server <id>] Displays BMP configuration.	
Syntax Description	VRF	Name default is "default VRF"

Default	N/A
Configuration Mode	config
History	3.7.1100
Example	

```
switch (config)# show ip bmp
-------------------------------------------------------------------------------------
ID          Admin State      Address            Port        Statistics Reporting Period
-------------------------------------------------------------------------------------
1           Active           1.1.1.1            11          20
2           Active           2.2.2.2            22          30
```

Related Commands	
Note	If no server ID is supplied, the command displays BMP configurations for all configured BMP servers under a VRF

15.4 Bidirectional Forwarding Detection (BFD) Infrastructure

Many protocols uses slow Hello mechanisms and failure is detection usually seconds after the problem occurs. The BFD goal is to provide low overhead short duration detection of failures between adjacent nodes and single mechanism that can be used for liveness detection over any media.

BFD session is established by the application that uses it. There is no discovery mechanism. E.g. in OSPF BFD session is established to neighbors that were discovered by OSPF hello protocol.

BFD supports multiple modes: one of them is Asynchronous.

In Asynchronous mode a system periodically sends BFD packets to verify connectivity. If a number of packets in a row are not received – the session is declared down.

A system can be passive or active. Active system initiates BDF sessions. Both systems can be active. (Only active mode is supported.)

> BFD is not yet supported for BGP unnumbered. Adjusting keepalive/hold timers may need to be considered to achieve faster convergence.

15.4.1 Session Establishment

A session begins with exchange of control packets. When bidirectional communication is achieved – a session becomes Up.

After session becomes up – control packet rate can be incremented.

Each side informs the neighbor in what intervals it is going to send BFD packets and what minimum interval it can receive BFD packets is.

Detection time is different in both directions and depends on negotiated parameters.

In Asynchronous mode—agreed transmit interval or remote system—max between local minimum rx time and last received min transmit time.

Detection time is equal to agreed transmit interval of remote system multiplied to multiplier received from remote system.

15.4.2 Interaction with Protocols

BFD session can be single-hop or multi-hop:

- Single hop session traverse between two adjacent IP neighbors. BFD control packet should be encapsulated in UDP with DPORT = 3784. SPORT should be in range 49152 to 65535. Same SPORT must be used for all control BFD packets for given session and is unique between different sessions. TTL value is 255.
- Multi-hop sessions traverse between to remote ip neighbors. Control packets are encapsulated in UDP with DPORT = 4784.

If different protocols want to establish a BFD session with the same remote system for same data plane – they should share BFD session.

IPv4 and IPv6 data protocols have different BFD sessions.

In OSPF Protocol neighbor discovery protocol establishes single hop BFD sessions. For OSPF when session fails – it tears down OSPF neighbor.

BFD session is established to BGP neighbor (single hop or multiple hop).

Single hop BFD session can be established for static route next hop.

15.4.3 BFD Commands

15.4.3.1 protocol bfd

	protocol bfd no protocol bfd Enables bfd on a system level The no form of the command removes bfd configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	config router bgp
History	3.6.4070
Example	`switch (config router bgp)# protocol bfd`
Related Commands	
Notes	The command returns an error if BFD is enabled in clients already running on the system (static routes or BGP of OSPF)

15.4.3.2 bfd shutdown

	bfd shutdown [vrf <vrf-name>] no bfd shutdown [vrf <vrf-name>] Disables bfd sessions but doesn't remove the configuration. if VRF is not given the command will be executed in active VRF.
Syntax Description	N/A
Default	N/A
Configuration Mode	config router bgp
History	3.6.4070
Example	`switch (config) # ip bfd shutdown`
Related Commands	
Notes	• The command “no ip bfd shutdown” or BFD interval parameters modification are affect traffic for all protocols; OSPF, BGP, static routes. The dynamic protocols (OSPF and BGP) restore the connection based on Hello protocol. • For static routes, please execute “no ip route static bfd <ip address>”

15.4.3.3 bfd interval

	bfd interval [vrf <vrf-name>] [transmit-rate] [min-rx] [multiplier] no bfd interval Sets the interval rates between BFD messages. The no form of the command removes bfd interval rates.	
Syntax Description	transmit-rate	Transfer time between two consecutive BFD messages, the actual time is negotiated between two systems Range: 50-60000 (msec)
	min_rx	Minimum time between neighbor messages, the actual time is negotiated between two systems Range: 50-60000 (msec)
	multiplier	Defines a time period to detect BFD failure Range: 3-50
Default	transmit-rate - 300 min-rx - 150 multiplier - 3	
Configuration Mode	config	
History	3.6.4070	
Example	`switch (config) # ip bfd interval transmit-rate 300 multiplier 3 min-rx 300 force`	
Related Commands		
Notes	The command is executed in the active VRF if a VRF is not specified	

15.4.3.4 ip ospf bfd

	ip ospf bfd no ip ospf bfd Enables BFD on the given interface for all OSPF neighbors on a number of active sessions. The no form of the command disables BFD on all OSPF neighbors.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config interface ethernet	
History	3.6.4070	
	3.6.4110	Added “no” form of the command
Example	`switch (config interface ethernet 1/2)# ip ospf bfd`	
Related Commands		
Notes	The command “ip ospf bfd” affects traffic, OSPF restores the connection based on Hello protocol	

15.4.3.5 ip route bfd

	ip route [vrf <vrf_name>] <prefix> <next_hop> bfd no ip route [vrf <vrf_name>] <prefix> <next_hop> bfd Configures static route with BFD enabled on a specified VRF. The no form of the commands removes the route.	
Syntax Description	vrf-name	VRF session name
	prefix	Subnet IP address
	next_hop	IP address of next hop
Default	N/A	
Configuration Mode	config	
History	3.6.4070	
	3.7.1100	Updated command syntax and Example
Example	`switch (config) # ip route vrf default 1.1.1.0/24 3.3.3.3 bfd`	
Related Commands		
Notes	When a session fails, all static routes pointing to the specified gateway are removed from the routing decision	

15.4.3.6 show ip route static

	show ip route [vrf [<vrf-name> | all]] static Displays static routing table of VRF instance.

Syntax Description	all	Displays routing tables for all VRF instances
	vrf	VRF name
Default	Default vrf	
Configuration Mode	Any command mode	
History	3.6.4070	
	3.7.1100	Update command syntax
Example	`switch (config) # show ip route vrf default static`	
Related Commands	ip route	
Notes	If no routing-context is specified, the “routing-context” VRF is automatically displayed	

15.4.3.7 show ip bfd neighbors

	show ip bfd [vrf <name> | all] neighbors [brief | <ip>] Displays BFD table of neighbor VRF instances.	
Syntax Description	all	Displays tables for all VRF instances
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4110	
Example		

```
switch (config) # show ip bgp neighbors 1000::1040

BGP neighbor: 1000::1040, remote AS: 100, link: external
  BGP version: 4, remote router ID: 2.1.1.1
  BGP State: ESTABLISHED
  Last read: 0:00:09:28, last write: 0:00:09:28, hold time is: 180, keepalive interval in seconds: 60
  BFD State: Up
  Configured hold time in seconds: 180, keepalive interval in seconds: 60
  Minimum holdtime from neighbor in seconds: 180

  Neighbor capabilities:
    Route refresh: advertise and received
    Graceful Restart Capability: advertise and received
    Address family IPv4 Unicast: advertise and received
    Address family IPv6 Unicast: n/a

  Message statistics:
    InQ depth is: 0
    OutQ depth is: 0
                            ----       -----
                             Sent       Rcvd
                            ----       -----
    Opens:                  1          1
    Notifications:          0          0
    Updates:                4          4
    Keepalives:             1587       1593
    Route Refresh:          0          0
    Total:                   1592       1598
  Default minimum time between advertisement runs in seconds: 30

  For address family: IPv4 Unicast
  BGP table version: 7
  Output queue size : 0

                                   ----       ----
                                    Sent       Rcvd
  Prefix activity:                  ----       ----
    Prefixes Current:               4          2
    Prefixes Total:                 4          2
    Implicit Withdraw:              0          0
    Explicit Withdraw:              0          0
    Used as bestpath:                 n/a     2
    Used as multipath:                n/a        n/a

                                   --------    -------
                                      Outbound    Inbound
  Local Policy Denied Prefixes:       --------    -------
    Total:                          0            0

  Connections established: 1; dropped: 1
  Last reset: 0:23:01:17, due to: 0 (0)
External BGP neighbor possible distance in hops: 1
Connection state is: ESTABLISHED
Local host: 1.1.1.1, Local port: 49616
Foreign host: 1000::1040, Foreign port: 179
```

Related Commands	
Notes	

15.5 Policy Rules

15.5.1 Route Map

Route maps define conditions for redistributing routes between routing protocols. A route map clause is identified by a name, filter type (permit or deny) and a sequence number. Clauses with the

same name are components of a single route map; the sequence number determines the order in which the clauses are compared to a route.

> Route maps can be used only for the BGP protocol.

> Route maps cannot be used for the commands “network” or “redistribute”.

15.5.2 Route Map Commands

- 15.5.1 Route Map
- 15.5.2 Route Map Commands
  - 15.5.2.1 route-map
  - 15.5.2.2 continue <sequence-number>
  - 15.5.2.3 abort
  - 15.5.2.4 match as-number
  - 15.5.2.5 match as-path
  - 15.5.2.6 match community-list
  - 15.5.2.7 match ip/ipv6 address
  - 15.5.2.8 match ip next-hop
  - 15.5.2.9 match metric
  - 15.5.2.10 set as-path prepend
  - 15.5.2.11 set community additive
  - 15.5.2.12 set community none
  - 15.5.2.13 set community delete
  - 15.5.2.14 set community-list
  - 15.5.2.15 set community-list additive
  - 15.5.2.16 set community-list delete
  - 15.5.2.17 set ip next-hop
  - 15.5.2.18 set local-preference
  - 15.5.2.19 set metric
  - 15.5.2.20 set origin
  - 15.5.2.21 set weight
  - 15.5.2.22 show route-map
  - 15.5.2.23 IP Prefix-List
    - 15.5.2.23.1 Configuring Prefix-List with Multiple Entries
  - 15.5.2.24 IP Prefix-List Commands
    - 15.5.2.24.1 ip prefix-list
    - 15.5.2.24.2 ip prefix-list bulk-mode
    - 15.5.2.24.3 ip prefix-list commit
    - 15.5.2.24.4 permit
    - 15.5.2.24.5 show ipv6 prefix-list

15.5.2.1 route-map

	route-map <map-name> [deny | permit] [sequence-number] no route-map <map-tag> {deny | permit} [<sequence-number>] Creates a route map that can be used for importing, exporting routes and applying local policies. The no form of the command deletes configured route maps.	
Syntax Description	name	Name of the route-map
	deny | permit	Configures the rule to be used
	sequence-number	Sequence number for a route-map specific record
Default	N/A	
Configuration Mode	config	
History	3.3.5006	
Example	`switch (config) # route-map mymap permit 1200` `switch (config route-map mymap permit 1200)#`	
Related Commands		
Notes	• All changes in a the route map configuration mode become pending until the end of the route-map session • If not configured, deny | permit is configured as permit • If not configured, sequence-number default value is 10	

15.5.2.2 continue <sequence-number>

	continue <sequence-number> no continue Enables additional route map evaluation of routes whose parameters meet the clause's matching criteria. The no form of the command removes this configuration from the route map clause.
Syntax Description	N/A
Default	N/A
Configuration Mode	config route map
History	3.3.5006

Example	`switch (config route-map mymap permit 10)# match as-number 40`
	`switch (config route-map mymap permit 10)# set weight 7`
	`switch (config route-map mymap permit 10)# continue 1200`
	`switch (config route-map mymap permit 10)# exit`

Related Commands	route-map <map-name> [deny | permit] [sequence-number]

Notes	• A clause typically contains a match (route-map) and a set (route-map) statement. The evaluation of routes whose settings are the same as match statement parameters normally end and the clause's set statement are applied to the route. Routes that match a clause containing a continue statement are evaluated against the clause specified by the continue statement. • When a route matches multiple route-map clauses, the filter action (deny or permit) is determined by the last clause that the route matches. The set statements in all clauses matching the route are applied to the route after the route map evaluation is complete. Multiple set statements are applied in the same order by which the route was evaluated against the clauses containing them. • Continue cannot be set to go back to a previous clause; <sequence-number> of the continue must always be higher than the current clause's sequence number.

15.5.2.3 abort

	abort Discards pending changes and returns to global configuration mode.
Syntax Description	N/A
Default	N/A
Configuration Mode	config route map
History	3.3.5006
Example	`switch (config route-map mymap permit 10)# abort`
Related Commands	
Notes	

15.5.2.4 match as-number

	match as-number <number> no match as-number Filters according to one of the AS numbers in the AS path of the route. The no form of the command removes this configuration from the route map clause.	
Syntax Description	number	Autonomous system number to check
Default	N/A	
Configuration Mode	config route map	
History	3.3.5006	
Example	`switch (config route-map mymap permit 10)# match as-number 40`	
Related Commands		
Notes	• When a clause contains multiple match commands, the permit or deny filter applies to a route only if its properties are equal to corresponding parameters in each match statement • When a route's properties do not equal the statement parameters, the route is evaluated against the next clause in the route map, as determined by sequence number • If all clauses fail to permit or deny the route, the route is denied	

15.5.2.5 match as-path

	match as-path <as-path-list name> no match as-path Creates a route map clause entry that matches the route‘s AS path using an as-path access-list. The no form of the command removes the match statement from the configuration mode route map clause.	
Syntax Description	number	Autonomous system number to check
Default	N/A	
Configuration Mode	config route map	
History	3.3.5006	
	3.6.3004	Added note
Example	`switch (config route-map mymap permit 10)# match as-path my-list`	
Related Commands		
Notes	• When a clause contains multiple match commands, the permit or deny filter applies to a route only if its properties are equal to corresponding parameters in each match statement • When a route’s properties do not equal the statement parameters, the route is evaluated against the next clause in the route map, as determined by sequence number • If all clauses fail to permit or deny the route, the route is denied • An as-path-list must already exist before a node is configured to use it	

15.5.2.6 match community-list

	match community <communities-list-name> exact-match no match community <communities-list-name> exact-match Creates a route map clause entry that specifies one route filtering condition. The no form of the command removes the match clause.	
Syntax Description	communities-list-name	A name of an IP community list
Default	N/A	
Configuration Mode	config route map	
History	3.3.5006	
Example	`switch (config route-map mymap permit 10)# match community-list COM_LIST exact-match`	
Related Commands		
Notes	• When a clause contains multiple match commands, the permit or deny filter applies to a route only if its properties are equal to corresponding parameters in each match statement. • When a route’s properties do not equal the statement parameters, the route is evaluated against the next clause in the route map, as determined by sequence number. • If all clauses fail to permit or deny the route, the route is denied.	

15.5.2.7 match ip/ipv6 address

	match ip address <prefix-list-name> no match ip address match ipv6 address <prefix-list-name> no match ipv6 address Filters according to IPv4/IPv6 prefix list. The no form of the command removes this configuration from the route map clause.	
Syntax Description	prefix-list-name	Prefix-list name
Default	N/A	
Configuration Mode	config route map	
History	3.3.5006	
Example	`switch (config route-map mymap permit 10)# match ip address listSmallRoutes`	
Related Commands		
Notes	• When a clause contains multiple match commands, the permit or deny filter applies to a route only if its properties are equal to corresponding parameters in each match statement • When a route's properties do not equal the statement parameters, the route is evaluated against the next clause in the route map, as determined by sequence number • If all clauses fail to permit or deny the route, the route is denied • The prefix-list-name should point to an existing IP prefix-list. If it is not found, no route is considered as a match for this clause.	

15.5.2.8 match ip next-hop

	match ip next-hop <ipv4/ipv6> no match ip next-hop Configures a route's entry next-hop match. The no form of the command removes a route-map's entry next-hop match.	
Syntax Description	ipv4/ipv6	Next hop IP address (e.g. 10.0.13.86)
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
	3.6.4070	Added support for IPv4 and IPv6
Example	`switch (config route-map mymap permit 10)# match ip next-hop 10.10.10.10`	
Related Commands		
Notes	• When a clause contains multiple match commands, the permit or deny filter applies to a route only if its properties are equal to corresponding parameters in each match statement • When a route's properties do not equal the statement parameters, the route is evaluated against the next clause in the route map, as determined by sequence number • If all clauses fail to permit or deny the route, the route is denied	

15.5.2.9 match metric

	match metric <value> no match metric Configures a route's entry metric match. The no form of the command removes a route-map's entry metric match.	
Syntax Description	value	Range: 1-2147483647.
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
	3.4.0000	Updated value range
Example	`switch (config route-map mymap permit 10)# match metric 10`	
Related Commands		
Notes	• When a clause contains multiple match commands, the permit or deny filter applies to a route only if its properties are equal to corresponding parameters in each match statement • When a route's properties do not equal the statement parameters, the route is evaluated against the next clause in the route map, as determined by sequence number • If all clauses fail to permit or deny the route, the route is denied	

15.5.2.10 set as-path prepend

	set as-path prepend <value$_1$> <value$_2$> ... <value$_n$> no set as-path prepend Modifies as-path on affected routes. The no form of the command removes the set statement from the route map.	
Syntax Description	value	BGP AS number that is prepended to as-path Range: 1-4294967295
Default	N/A	
Configuration Mode	config route map	
History	3.4.0000	
Example	`switch (config route-map mymap permit 10)# set as-path prepend 5 10`	
Related Commands		
Notes		

15.5.2.11 set community additive

	set community <list-of-communities> additive no set community <list-of-communities> additive Adds the matching communities. The no form of the command removes the set statement from the clause.

Syntax Description	list-of-communities	List of standard communities: • <aa:nn> • <number> • internet • local-AS • no-advertise • no-export
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
Example	`switch (config route-map mymap permit 10)# set community none`	
Related Commands		
Notes		

15.5.2.12 set community none

	set community none no set community none Sets the community attribute of a distributed route to be empty. The no form of the command removes the set statement from the clause.
Default	N/A
Configuration Mode	config route map
History	3.3.5200
Example	`switch (config route-map mymap permit 10)# set community none`
Related Commands	
Notes	

15.5.2.13 set community delete

	set community <list of communities> delete no set community <list of communities> delete Deletes matching communities. The no form of the command removes the set statement from the clause.	
Syntax Description	list of communities	List of standard communities: • <aa:nn> • <number> • internet • local-AS • no-advertise • no-export
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
Example	`switch (config route-map test_route_map permit 10) # set community 400:1 delete`	

Related Commands	
Notes	

15.5.2.14 set community-list

	set community-list <community-list-name> no set community <list of communities> Configures a named standard community list. The no form of the command removes the set statement from the clause.	
Syntax Description	<community-list-name>	Name of community list
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
Example	`switch (config route-map mymap permit 10 )# set community internet 1:3 additive`	
Related Commands		
Notes	A community-list must already exist before a node is configured to use it	

15.5.2.15 set community-list additive

	set community-list <community-list-name> additive no set community <list of communities> additive Adds to existing communities using the communities found in the community list. The no form of the command removes the set statement from the clause.	
Syntax Description	<community-list-name>	Name of community list
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
Example	`switch (config route-map mymap permit 10)# set community-list mycommunity additive`	
Related Commands		
Notes		

15.5.2.16 set community-list delete

	set community-list <community-list-name> delete no set community-list Deletes the matching community list permit entries from the route community list. The no form of the command removes the set statement from the clause.	
Syntax Description	community-list-name	Name of community list
Default	N/A	
Configuration Mode	config route map	

History	3.3.5200
Example	`switch (config route-map mymap permit 10)# set community-list mycommunity delete`
Related Commands	
Notes	

15.5.2.17 set ip next-hop

	set ip next-hop <ipv4/ipv6> no set ip next-hop Configures a route's entry next-hop parameter. The no form of the command removes a route-map's entry next-hop setting.	
Syntax Description	ipv4/ipv6	Route next-hop IP (e.g. 10.0.13.86)
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
	3.6.4070	Added support for IPv4 and IPv6
Example	`switch (config route-map mymap permit 10)# set ip next-hop 10.10.10.10`	
Related Commands		
Notes		

15.5.2.18 set local-preference

	set local-preference <value> no set local-preference Configures a route's entry local-preference parameter. The no form of the command removes a route-map's entry local-pref setting.	
Syntax Description	value	Route local-pref Range: 1-2147483648
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
Example	`switch (config route-map mymap permit 10)# set local-preference 10`	
Related Commands		
Notes		

15.5.2.19 set metric

	set metric <value> no set metric Configures a route's entry metric parameter. The no form of the command removes a route-map's entry metric setting.

Syntax Description	value	Route metric Range: 1-2147483647
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
Example	`switch (config route-map mymap permit 10)# set metric 10`	
Related Commands		
Notes		

15.5.2.20 set origin

	set origin <egp | igp | incomplete> no set origin Configures a route's entry origin parameter. The no form of the command removes a route-map's entry origin setting.	
Syntax Description	egp	Set a route's entry origin parameter to external.
	igp	Set a route's entry origin parameter to internal.
	incomplete	Set a route's entry origin parameter to incomplete.
Default	N/A	
Configuration Mode	config route map	
History	3.3.5200	
Example	`switch (config route-map mymap permit 10)# set origin egp`	
Related Commands		
Notes		

15.5.2.21 set weight

	set weight <number> no set weight Configures modifications to redistributed routes. The no form of the command removes this configuration from the route map clause.	
Syntax Description	number	Value of the weight to set Range: 1-65535
Default	N/A	
Configuration Mode	config route map	
History	3.3.5006	
	3.4.0000	Updated parameter range
Example	`switch (config route-map mymap permit 10)# set weight 7`	
Related Commands	route-map <map-name> [deny | permit] [sequence-number]	
Notes		

15.5.2.22 show route-map

	show route-map [<name>] Displays route map configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.5006

Example	`switch (config)# show route-map mymap`
	`route-map mymap, permit, sequence 10`
	`  Match clauses:`
	`  as-number 40`
	`  Set clauses:`
	`  weight 7`
	`route-map mymap, permit, sequence 1200`
	`  Set clauses:`
	`  weight 11`

Related Commands	
Notes	

15.5.2.23 IP Prefix-List

IP prefix-lists are used to match two components of IP packets or an IP route. Prefix-list is a list of entries that include an IP network address and a bit mask (Range: 1 to 32 and should match the input IP network address).

15.5.2.23.1 Configuring Prefix-List with Multiple Entries

To create a new prefix-list with a large number of entries (50K for IPv4 or 25K for IPv6), use "configuration text fetch" to fetch a predefined prefix-list configuration file and then apply it as a whole.
In order to edit an existing prefix-list, the maximum entries that can be updated every time is 1K at most. An update operation of more than 1K entries can be achieved by doing this multiple times.

Configuration fetch example where fetch “prefix-list-001”:

```
switch (config) # configuration text fetch ?

 <download
 URL>
 http, https, ftp, tftp, scp and sftp are supported.  e.g.
 scp://username[:password]@hostname/path/filename
```

Apply:

```
switch (config) # configuration text file prefix-list-001 apply verbose
All commands succeeded.

Transcript of all commands executed:

------------ Begin transcript ------------
    Onyx-Demo (config) # ip prefix-list prefix-list-001
    Onyx-Demo (config) # seq 1 permit 200.1.1.0 eq 24
    Onyx-Demo (config) # seq 2 permit 1.1.1.2 eq  32
    Onyx-Demo (config) # seq 3 permit 1.1.1.3 eq  32
    Onyx-Demo (config) # seq 4 permit 1.1.1.4 eq  32
    Onyx-Demo (config) # seq 5 permit 1.1.1.5 eq  32
    Onyx-Demo (config) # seq 6 permit 1.1.1.6 eq  32
```

```
Onyx-Demo (config) # seq 7 permit 1.1.1.7 eq  32
Onyx-Demo (config) # seq 8 permit 1.1.1.8 eq  32
Onyx-Demo (config) # exit
 ------------ End transcript ------------
```

15.5.2.24 IP Prefix-List Commands

15.5.2.24.1 ip prefix-list

	ip prefix-list <list-name> [seq <number>] no ip prefix-list <list-name> [seq <number>] ipv6 prefix-list <list-name> [seq <number>] no ipv6 prefix-list <list-name> [seq <number>] Configures or updates the IPv4 or IPv6 prefix-list in context mode. The no form of the command deletes the prefix-list or a prefix-list entry.	
Syntax Description	list-name	String
	seq <number>	Sequence number assigned to entry Range: 0-4294967295 Default value: 10
Default	N/A	
Configuration Mode	config	
History	3.3.5200	
	3.6.4070	Added support for IPv6
	3.8.2100	Updated maximum sequence value. Reorganized the command into ip prefix-list command and sub-commands.
Example	switch (config) # ip prefix-list list-name switch (config ip prefix-list list-name) # deny 1.1.1.0 /24 switch (config ip prefix-list list-name) # deny 1.1.2.0 /24 switch (config ip prefix-list list-name) # exit switch (config) # switch (config) # show ip prefix-list list-name prefix-list list-name: count: 2, range entries: 0, sequences: 10 - 20 Configuration: seq 10 deny 1.1.1.0 /24 eq 24 seq 20 deny 1.1.2.0 /24 eq 24	
Related Commands	route-table prefix-list show ip bgp vrf address-family	
Notes	The maximum entries for IPv4 prefix-list is 50K and for IPv6 is 25K.	

15.5.2.24.2 ip prefix-list bulk-mode

	ip prefix-list <list-name> bulk-mode no ip prefix-list <list-name> bulk-mode Enables bulk-mode for a given prefix-list. Disables bulk-mode for a given prefix-list.	
Syntax Description	list-name	String
Default	N/A	

Configuration Mode	config
History	3.9.1900
Example	switch (config) # ip prefix-list list-name switch (config) # ip prefix-list list-name bulk-mode # bulk-mode will be enabled for the prefix-list switch (config) # ip prefix-list list-name seq 10 permit 20.20.20.20 /32 eq 32 switch (config) # ip prefix-list list-name seq 20 deny 21.21.21.21 /32 eq 32 switch (config) # ip prefix-list list-name commit # bulk setting of rules applied to Onyx, and bulk-mode for this prefix list is cleared.
Related Commands	
Notes	• In case of bulk-mode enabled, the prefix list rule configuration will be cached in CLI until 'commit' command is issued. Otherwise, the rule configuration will be applied immediately. • To apply prefix list configuration in bulk-mode will improve performance greatly in case of a very large prefix list (50K and up). The bulk mode is enabled by default if prefix list rules are configured under CLI prefix mode. When 'exit' is issued to quit from the CLI prefix mode, CLI will aggregate all the rule configuration and apply the bulk setting to the system.

15.5.2.24.3 ip prefix-list commit

	ip prefix-list <list-name> commit If bulk-mode is enabled for the prefix list, then commit the whole prefix-list configuration and reset bulk mode (otherwise, nothing will happen).	
Syntax Description	list-name	String
Default	N/A	
Configuration Mode	config	
History	3.9.1900	
Example	switch (config) # ip prefix-list list-name commit	
Related Commands		
Notes		

15.5.2.24.4 permit

	[seq <number>] <permit|deny> <ipv4_address|ipv6_address> <mask> [eq <length> | le <length> | ge <length> [le <length>]] Configures IPv4 or IPv6 permit/deny clauses.	
Syntax Description	permit | deny	Configures the prefixes to be used
	ipv4_address	IPv4 address
	Ipv6_address	IPv6 address
	eq | ge | le <mask>	• eq—equal to a specified prefix length • ge—greater than or equal to a specified prefix length • le—less than or equal to a specified prefix length
Default	N/A	
Configuration Mode	config	
History	3.8.2100	

Example	switch (config) # ip prefix-list list-name switch (config ip prefix-list list-name) # deny 1.1.1.0 /24 switch (config ip prefix-list list-name) # deny 1.1.2.0 /24 switch (config ip prefix-list list-name) # exit switch (config) # switch (config) # show ip prefix-list list-name prefix-list list-name: count: 2, range entries: 0, sequences: 10 - 20 Configuration: seq 10 deny 1.1.1.0 /24 eq 24 seq 20 deny 1.1.2.0 /24 eq 24
Related Commands	route-table prefix-list show ip bgp vrf address-family
Notes	

15.5.2.24.5 show ipv6 prefix-list

	show ipv6 prefix-list [<name>] Displays IPv6 prefix-lists.	
Syntax Description	name	Displays a specific prefix-list
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.6.4070	Added support for IPv6
Example	switch (config)# show ipv6 prefix-list prefix-list: a-list count: 1, range entries: 1, sequences: 10 - 10 seq 10 permit 2001::0 /64 ge eq 32 (hit count: 0, refcount: 0)	
Related Commands		
Notes		

15.6 VRRP

The Virtual Router Redundancy Protocol (VRRP) is a computer networking protocol that provides for automatic assignment of available IP routers to participating hosts. This increases the availability and reliability of routing paths via automatic default gateway selections on an IP subnetwork.

The protocol achieves this by creating virtual routers, which are an abstract representation of multiple routers (that is, a master and backup routers, acting as a group). The default gateway of a participating host is assigned to the virtual router instead of a physical router. If the physical router that is routing packets on behalf of the virtual router fails, another physical router is selected to automatically replace it. The physical router that is forwarding packets at any given time is called the master router.

VRRP provides information on the state of a router, not the routes processed and exchanged by that router. Each VRRP instance is limited, in scope, to a single subnet. It does not advertise IP routes beyond that subnet or affect the routing table in any way.

Routers have a priority of between 1-255 and the router with the highest priority becomes the master. The configurable priority value ranges from 1-254, the router which owns the interface IP address as one of its associated IP addresses has the priority value 255. When a planned withdrawal of a master router is to take place, its priority can be lowered, which means a backup router will preempt the master router status rather than having to wait for the hold time to expire.NVIDIA Onyx supports IPv4 in VRRP version 2, and IPv6 in VRRP version 3.

15.6.1 Load Balancing

To create load balancing between routers participating in the same VR, it is recommended to create 2 (or more) VRs. Each router will be a master in one of the VRs, and a backup to the other VR(s). A group of hosts should be configured with Router 1's virtual address as the default gateway, while the second group should be configured with Router 2's virtual address.

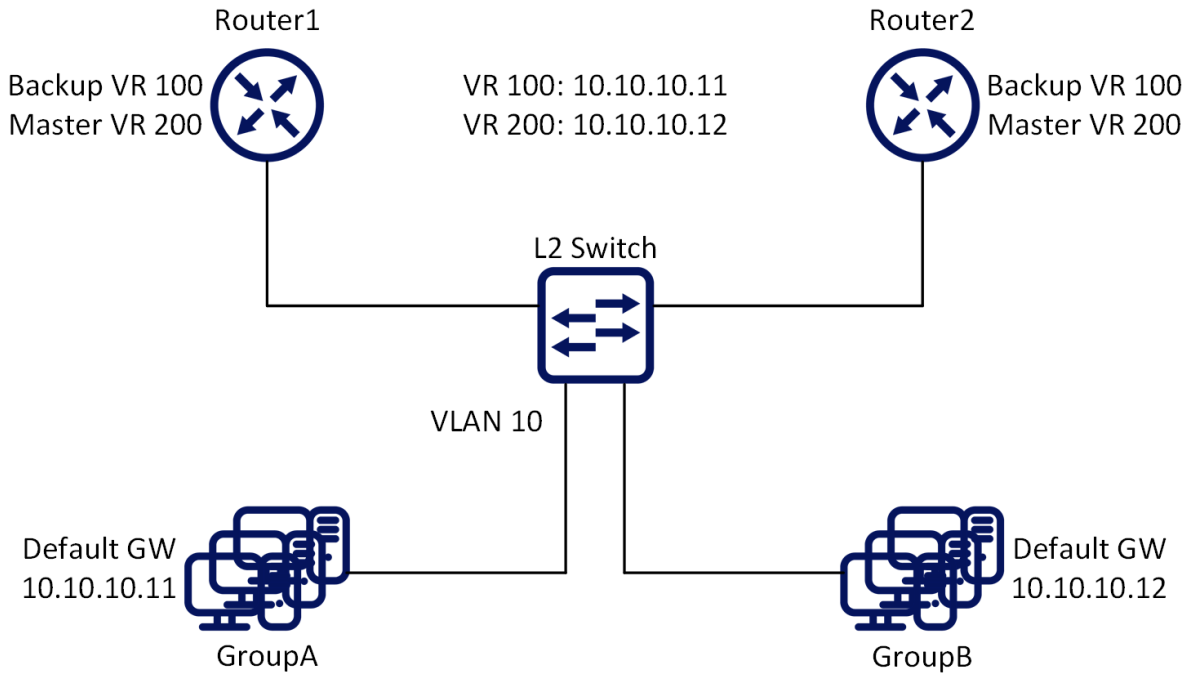

15.6.2 Configuring VRRP

The Virtual Router Redundancy Protocol (VRRP) is a computer networking protocol that provides for automatic assignment of available IP routers to participating hosts. This increases the availability and reliability of routing paths via automatic default gateway selections on an IP subnetwork.The protocol achieves this by creating virtual routers, which are an abstract representation of multiple routers (that is, a master and backup routers, acting as a group). The default gateway of a participating host is assigned to the virtual router instead of a physical router. If the physical router that is routing packets on behalf of the virtual router fails, another physical router is selected to automatically replace it. The physical router that is forwarding packets at any given time is called the master router.VRRP provides information on the state of a router, not the routes processed and exchanged by that router. Each VRRP instance is limited, in scope, to a single subnet. It does not advertise IP routes beyond that subnet or affect the routing table in any way.Routers have a priority of between 1-255 and the router with the highest priority becomes the master. The configurable priority value ranges from 1-254, the router which owns the interface IP address as one of its

associated IP addresses has the priority value 255. When a planned withdrawal of a master router is to take place, its priority can be lowered, which means a backup router will preempt the master router status rather than having to wait for the hold time to expire.

15.6.2.1 Preconditions

1. Enable IP routing functionality. Run:

```
switch (config)# ip routing
```

2. Enable the desired VLAN. Run:

```
switch (config)# vlan 20
```

> The VLAN cannot be the same one configured for the MLAG IPL, if MLAG is used.

3. Add this VLAN to the desired interface. Run:

```
switch (config)# interface ethernet 1/1
switch (config interface ethernet 1/1)# switchport access vlan 20
```

4. Create a VLAN interface. Run:

```
switch (config)# interface vlan 20
```

5. Apply IP address to the VLAN interface.
   a. For IPv4, do the following.
      On one of the switches, run:

```
switch (config interface vlan 20)# ip address 20.20.20.20 /24
```

      On the other switch, run:

```
switch (config interface vlan 20)# ip address 20.20.20.30 /24
```

   b. For IPv6, apply IPv6 address to the VLAN interface.

      On one of the switches, run:

```
switch (config interface vlan 20) # ipv6 address 2001::20 /64
```

      On the other switch, run:

```
switch (config interface vlan 20) # ipv6 address 2001::30 /64
```

6. Enable the interface. Run:

```
switch (config interface vlan 20)# no shutdown
```

15.6.2.2 Configuring VRRP

1. Enable VRRP protocol globally. Run:

```
switch (config)# protocol vrrp
```

2. Create a virtual router group for an IP interface. Up to 255 VRRP IDs are supported. Run:

```
switch (config interface vlan 20)# vrrp 100
```

3. Set the VIP address.
   a. For IPv4, run:

```
switch (config interface vlan 20 vrrp 100)# address 20.20.20.40
```

   b. For IPv6, run:

```
switch (config interface vlan 20 vrrp 100) # address 2001::40
```

4. Influence the election of the master in the VR cluster make sure that the priority of the desired master is the highest. Note that the higher IP address is selected in case the priority of the routers in the VR are the same. Select the priority. Run:

```
switch (config interface vlan 20 vrrp 100)# priority 200
```

5. The advertisement interval should be the same for all the routers within the VR. Modify the interval. Run:

```
switch (config interface vlan 20 vrrp 100)# advertisement-interval 2
```

6. The authentication text should be the same for all the routers within the VR. Configure the authentication text. Run:

```
switch (config interface vlan 20 vrrp 100)# authentication text my-password
```

   This option is not supported in VRRP IPv6.

7. Use the preempt command to enable a high-priority backup virtual router to preempt the low-priority master virtual router. Run:

```
switch (config interface vlan 20 vrrp 100)# preempt
```

8. Disable VRRP. Run:

```
switch (config interface vlan 20 vrrp 100)# shutdown
```

   The configuration will not be deleted, only the VRRP state machine will be stopped.

15.6.2.3 Verifying VRRP

1. Display VRRP brief status. Run:

```
switch (config) # show vrrp
Interface  VR         Admin State      Priority   Adv-Intvl   Preempt    State      VR IP addr
-----------------------------------------------------------------------------------------------------------
Vlan20     100        Enabled           100        1          Enabled    Master     20.20.20.40
Vlan20     100        Enabled           100        1          Enabled    Master     2001::40
```

2. Display VRRP detailed status. Run:

```
switch (config) # show vrrp detail
VRRP Admin State: Enabled

Vlan20 - Vrrp 100:
  Instance Admin State         : Enabled
  State                        : Master
  State v6                     : Master
  Virtual IP Address           : 20.20.20.40
  Virtual IPv6 Address         : 2001::40
  Priority                     : 100
  Advertisement interval(sec)  : 1
  Preemption                   : Enabled
  Virtual MAC Address          : 00:00:5e:00:01:64
  Primary IP Address           : 20.20.20.20
  Master router                : 20.20.20.20
  Virtual MAC Address v6       : 00:00:5e:00:02:64
  Primary IP Address v6        : ::
  Master router v6             : 2001::20
  Master priority              : 100
  Master advertisement interval: 1
```

3. Display VRRP statistic counters. Run:

```
switch (config) # show vrrp statistics
Invalid packets:                    0
Too short:                          0
Transitions to Master:              1
Total received:                     0
Bad TTL:                            0
Failed authentication:              0
Unknown authentication:             0
Conflicting authentication:         0
Conflicting Advertise time:         0
Conflicting Addresses:              0
Received with zero priority:        0
Sent with zero priority:            0
Invalid packets v6:                 0
Too short v6:                       0
Transitions to Master v6:           1
Total received v6:                  0
Bad TTL v6:                         0
Conflicting Advertise time v6:      0
Conflicting Addresses v6:           0
Received with zero priority v6:     0
Sent with zero priority v6:         0
```

15.6.3 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community post:

- HowTo Configure VRRP on Ethernet Switches

15.6.4 VRRP Commands

- 15.6.1 Load Balancing
- 15.6.2 Configuring VRRP
  - 15.6.2.1 Preconditions
  - 15.6.2.2 Configuring VRRP

- 15.6.2.3 Verifying VRRP
- 15.6.3 Additional Reading and Use Cases
- 15.6.4 VRRP Commands
  - 15.6.4.1 protocol vrrp
  - 15.6.4.2 clear vrrp statistics
  - 15.6.4.3 vrrp
  - 15.6.4.4 address
  - 15.6.4.5 shutdown
  - 15.6.4.6 priority
  - 15.6.4.7 preempt
  - 15.6.4.8 authentication text
  - 15.6.4.9 advertisement-interval
  - 15.6.4.10 show vrrp
  - 15.6.4.11 show vrrp detail
  - 15.6.4.12 show vrrp statistics

15.6.4.1 protocol vrrp

	protocol vrrp no protocol vrrp Enables VRRP globally and unhides VRRP related commands. The no form of the command deletes all the VRRP configuration and hides VRRP related commands.
Syntax Description	N/A
Default	no protocol vrrp
Configuration Mode	config
History	3.3.4500
Example	`switch (config)# protocol vrrp`
Related Commands	
Notes	

15.6.4.2 clear vrrp statistics

	clear vrrp statistics Clears VRRP statistics.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.3.4500
Example	`switch (config)# clear vrrp statistics`
Related Commands	
Notes	

15.6.4.3 vrrp

	vrrp <number> no vrrp <number> Creates a virtual router group on this interface and enters a new configuration mode. The no form of the command deletes the VRRP instance and the related configuration.	
Syntax Description	number	A VRRP instance number Range: 1-255
Default	N/A	
Configuration Mode	config interface vlan	
History	3.3.4500	
	3.6.8100	Updated parameter range
	3.7.1100	Updated Syntax and notes
Example	`switch (config interface vlan 10)#` `switch (config interface vlan 10 vrrp 10)#`	
Related Commands		
Notes	A maximum total of 64 VRRP instances are supported per switch system.	

15.6.4.4 address

	address <ip-address> [secondary] no address [<ip-address> [secondary]] Sets virtual router IP address (primary and secondary). The no form of the command deletes the IP address from the VRRP interface.	
Syntax Description	ip-address	The virtual IP address
	secondary	A secondary IP address for the virtual router
Default	N/A	
Configuration Mode	config vrrp interface	
History	3.3.4500	
	3.9.1000	Added support IPv6 address

Example	`switch (config vrrp 100)# address 10.10.10.10`
	`switch (config vrrp 100)# address 10.10.10.11 secondary`
	`switch (config vrrp 100)# address 10.10.10.12 secondary`
	
	`switch (config vrrp 100)# address 2001::40`
	`switch (config vrrp 100)# address 2001::41 secondary`	

Related Commands		

Notes	• The virtual address can be either from the interface's primary or secondary subnet • This command is the enabler of the protocol. Therefore, set all the protocol parameters initially and only then set the ip-address. • There are up to 20 IP addresses associated with the VRRP instance. One primary and up to 19 secondary ip-addresses. • If the configured IP address is the same as the interface IP address, this switch automatically owns the IP address (priority 255) • For IPv6, the OS will auto-generate link-local virtual IP. Up to 19 IPv6 addresses are allowed to be associated with the VRRP instance—one primary address and up to 18 secondary addresses. IPv4 and IPv6 addresses are allowed to be configured on the same VRRP instance.

15.6.4.5 shutdown

	shutdown no shutdown Disables the virtual router. The no form of the command enables the virtual router (stops the VRRP state machine).
Syntax Description	N/A
Default	Enabled (no shutdown)
Configuration Mode	config vrrp interface
History	3.3.4500
Example	`switch (config vrrp 100)# shutdown`
Related Commands	
Notes	

15.6.4.6 priority

	priority <level> no priority Sets the priority of the virtual router. The no form of the command resets the priority to its default.	
Syntax Description	level	The virtual router priority level Range: 1-254
Default	100	
Configuration Mode	config vrrp interface	
History	3.3.4500	
Example	`switch (config vrrp 100)# priority 200`	
Related Commands		
Notes	• The higher IP address is selected as master if the priority of the routers in the VR are the same • To influence the election of the master in the VR cluster make sure that the priority of the desired master is the higher	

15.6.4.7 preempt

	preempt no preempt Sets virtual router preemption mode. The no form of the command disables the virtual router preemption.
Syntax Description	N/A
Default	Enabled (preempt)
Configuration Mode	config vrrp interface
History	3.3.4500
Example	`switch (config vrrp 100)# preempt`
Related Commands	
Notes	To set this router as backup for the current virtual router master, preempt must be enabled.

15.6.4.8 authentication text

	authentication text <password> no authentication text Sets virtual router authentication password and enables authentication. The no form of the command disables the authentication mechanism.	
Syntax Description	password	The virtual router authentication password
Default	Disabled	
Configuration Mode	config vrrp interface	
History	3.3.4500	
	3.9.1000	Updated notes
Example	`switch (config vrrp 100)# authentication text mypassword`	
Related Commands		
Notes	• The password string must be up to 8 alphanumeric characters • This option is not supported in VRRP IPv6 instance	

15.6.4.9 advertisement-interval

	advertisement-interval <seconds> no advertisement-interval Sets the virtual router advertisement-interval. The no form of the command resets the parameter to its default.	
Syntax Description	seconds	The virtual router advertisement-interval in seconds Range: 1-255
Default	1	
Configuration Mode	config vrrp interface	
History	3.3.4500	

Example	`switch (config vrrp 100)# advertisement-interval 10`
Related Commands	
Notes	

15.6.4.10 show vrrp

	show vrrp [interface <type> <number>] [vr <id>] Displays VRRP brief configuration and status.	
Syntax Description	interface <type> <number>	Filters the output to a specific interface type and number
	vr <id>	Filters the output to a specific virtual router Range: 1-10
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4500	
	3.9.1000	Added support for VRRP IPv6 instance

Example	`switch (config) # show vrrp`
	`Interface VR Admin State Priority Adv-Intvl Preempt State VR IP addr`
	`---------------------------------------------------------------------------`
	`Vlan20 100 Enabled 100 1 Enabled Master 20.20.20.40`
	`Vlan20 100 Enabled 100 1 Enabled Master 2001::40`	

Related Commands		
Notes		

15.6.4.11 show vrrp detail

	show vrrp detail [interface <type> <number>] [vr <id>] Displays detailed VRRP configuration and status.	
Syntax Description	interface <type> <number>	Filters the output to a specific interface type and number
	vr <id>	Filters the output to a specific virtual router Range: 1-255
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4500	
	3.6.5000	Updated example
	3.6.8008	Updated example
	3.9.1000	Added support for VRRP IPv6 instance

Example	switch (config) # show vrrp detail VRRP Admin State: Enabled Vlan20 - Vrrp 100: Instance Admin State : Enabled State : Master State v6 : Master Virtual IP Address : 20.20.20.40 Virtual IPv6 Address : 2001::40 Priority : 100 Advertisement interval(sec) : 1 Preemption : Enabled Virtual MAC Address : 00:00:5e:00:01:64 Primary IP Address : 20.20.20.20 Master router : 20.20.20.20 Virtual MAC Address v6 : 00:00:5e:00:02:64 Primary IP Address v6 : :: Master router v6 : fe80::ba59:9fff:fea6:6988 Master priority : 100 Master advertisement interval: 1 Associated IP Addresses: 20.20.20.41 Associated IPv6 Addresses: 2001::41
Related Commands	
Notes	

15.6.4.12 show vrrp statistics

	show vrrp statistics [interface <type <number>] [vr <id>] [all] Displays VRRP counters.	
Syntax Description	interface <type> <number>	Filters the output to a specific interface type and number
	vr <id>	Filters the output to a specific virtual router Range: 1-255
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4500	
	3.6.5000	Updated example
	3.9.1000	Updated example

Example	switch (config) # show vrrp statistics Invalid packets: 0 Too short: 0 Transitions to Master: 0 Total received: 0 Bad TTL: 0 Failed authentication: 0 Unknown authentication: 0 Conflicting authentication: 0 Conflicting Advertise time: 0 Conflicting Addresses: 0 Received with zero priority: 0 Sent with zero priority: 0 Invalid packets v6: 0 Too short v6: 0 Transitions to Master v6: 0 Total received v6: 0 Bad TTL v6: 0 Conflicting Advertise time v6: 0 Conflicting Addresses v6: 0 Received with zero priority v6: 0 Sent with zero priority v6: 0
Related Commands	
Notes	

15.7 MAGP

Multi-active gateway protocol (MAGP) is aimed to solve the default gateway problem when a host is connected to a set of switch routers (SRs) via MLAG.

The network functionality in that case requires that each SR is an active default gateway router to the host, thus reducing hops between the SRs and directly forwarding IP traffic to the L3 cloud regardless which SR traffic comes through.

15.7.1 Configuring MAGP

15.7.1.1 Prerequisites

1. Enable IP routing functionality. Run:

```
switch (config)# ip routing
```

2. Enable the desired VLAN. Run:

```
switch (config)# vlan 20
switch (config vlan 20)#
```

> The VLAN cannot be the same one configured for the MLAG IPL, if MLAG is used.

3. Add this VLAN to the desired interface. Run:

```
switch (config)# interface ethernet 1/1
switch (config interface ethernet 1/1)# switchport access vlan 20
```

4. Create a VLAN interface. Run:

```
switch (config)# interface vlan 20
switch (config interface vlan 20)#
```

5. Set an IP address to the VLAN interface.
   a. For IPv4, run:

```
switch (config interface vlan 20)# ip address 11.11.11.11 /8
```

   b. For IPv6, run:

```
switch (config interface vlan 20)# ip address 2001::11 /64
```

6. Enable the interface.

```
switch (config interface vlan 20)# no shutdown
```

15.7.1.2 Configuring MAGP

1. Enable MAGP protocol globally. Run:

```
switch (config)# protocol magp
```

2. Create a virtual router group for an IP interface. Run:

```
switch (config interface vlan 20)# magp 100
```

   Up to 255 MAGP IDs are supported.

3. Set a virtual router primary IP address.
   a. For IPv4, run:

```
switch (config interface vlan 20 magp 100)# ip virtual-router address 11.11.11.254
```

   b. For IPv6, run:

```
switch (config interface vlan 20 magp 100)# ip virtual-router address 2001::254
```

   Only a virtual IP from the primary subnet can be configured for MAGP.

4. Set a virtual router primary MAC address. Run:

```
switch (config interface vlan 20 magp 100)# ip virtual-router mac-address aa:bb:cc:dd:ee:ff
```

   To obtain the virtual router's MAC address, please run the command "show vrrp detail".

15.7.1.3 Verifying MAGP

To verify the MAGP configuration, run:

```
switch (config) # show magp

MAGP 100:
  Interface vlan: 20
  Admin state   : Enabled
  State         : Master
  Virtual IP    : 11.11.11.254
  V6 State      : Master
  Virtual IPv6  : 2001::254
  Virtual MAC   : aa:bb:cc:dd:ee:ff
```

This output is to be expected in both MAGP switches.

15.7.2 Useful Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community post:

- HowTo Configure MLAG+MAGP: Running Config Example
- HowTo Configure MAGP

15.7.3 MAGP Commands

- 15.7.1 Configuring MAGP
  - 15.7.1.1 Prerequisites
  - 15.7.1.2 Configuring MAGP
  - 15.7.1.3 Verifying MAGP
- 15.7.2 Useful Reading and Use Cases
- 15.7.3 MAGP Commands
  - 15.7.3.1 protocol magp
  - 15.7.3.2 magp
  - 15.7.3.3 shutdown
  - 15.7.3.4 ip virtual-router address
  - 15.7.3.5 ip virtual-router mac-address
  - 15.7.3.6 ip virtual-router mac-address <address>
  - 15.7.3.7 show magp
  - 15.7.3.8 show magp interface vlan

15.7.3.1 protocol magp

	protocol magp no protocol magp Enables MAGP globally and unhides MAGP commands. The no form of the command deletes all the MAGP configuration and hides MAGP commands.
Syntax Description	N/A

Default	Disabled
Configuration Mode	config
History	3.3.4500
Example	`switch (config)# protocol magp`
Related Commands	
Notes	IP routing must be enabled to enable MAGP.

15.7.3.2 magp

	magp <instance> no magp <instance> Creates an MAGP instance on this interface and enters a new configuration mode. The no form of the command deletes the MAGP instance.	
Syntax Description	instance	MAGP instance number Range: 1-255
Default	Disabled	
Configuration Mode	config interface vlan	
History	3.3.4500	
	3.7.1100	Updated notes
Example	`switch (config interface vlan 20)# magp 100` `switch (config interface vlan 20 magp 100)#`	
Related Commands		
Notes	• Only one MAGP instance can be created on an interface • Different interfaces cannot share an MAGP instance • MAGP and VRRP are mutually exclusive • A maximum total of 64 MAGP instances are supported per switch system	

15.7.3.3 shutdown

	shutdown no shutdown Enables MAGP instance. The no form of the command disables the MAGP instance.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface vlan magp
History	3.3.4500
Example	`switch (config interface vlan 10 magp 1)# shutdown`
Related Commands	
Notes	

15.7.3.4 ip virtual-router address

	ip virtual-router address <ip-address> [secondary] no ip virtual-router address <ip-address> [secondary] Sets MAGP virtual IP address. The no form of the command resets this parameter to its default.	
Syntax Description	ip-address	The virtual router IP address
	secondary	Adds secondary virtual router address
Default	N/A	
Configuration Mode	config interface vlan magp	
History	3.3.4500	
	3.6.8100	Added “secondary” parameter
	3.9.1000	Added support for MAGP IPv6 instance
Example	`switch (config interface vlan 10 magp 1)# ip virtual-router address 10.10.10.10` `switch (config interface vlan 20 magp 100) # ip virtual-router address 2001::254`	
Related Commands		
Notes	• The MAGP virtual IP address must be different from the interface IP address • In a single MAGP instance, IPv4 and IPv6 addresses are both allowed	

15.7.3.5 ip virtual-router mac-address

	ip virtual-router mac-address <mac-address> no ip virtual-router mac-address Sets MAGP virtual MAC address. The no form of the command resets the MAC address to its default.	
Syntax Description	mac-address	MAC address (format: aa:bb:cc:dd:ee:ff)
Default	00:00:5E:00:01-<magp instance>	
Configuration Mode	config interface vlan magp	
History	3.3.4500	
	3.9.1000	Added note about MAGP IPv6
	3.9.3000	Updated MAC address to be lowercase
Example	`switch (config interface vlan 10 magp 1)# ip virtual-router mac-address aa:bb:cc:dd:ee:ff`	
Related Commands		
Notes	• If not defined, "ip virtual-router mac-address <address>" address is used • If the "ip virtual-router mac-address <address>" is not defined, the default is used • In a single MAGP instance, IPv4 and IPv6 use a single virtual MAC	

15.7.3.6 ip virtual-router mac-address <address>

	ip virtual-router mac-address <address> Sets a global virtual router MAC address.

Syntax Description	address	MAC address (format: aa:bb:cc:dd:ee:ff)
Default	N/A	
Configuration Mode	config	
History	3.9.0500	
Example	`switch (config)# ip virtual-router mac-address 00:00:5E:00:11:22`	
Related Commands	show ip routing	
Notes	• The system can have only one MAC address • If this address is in use, it cannot be changed or removed	

15.7.3.7 show magp

	show magp [<instance>] Displays the MAGP configuration.	
Syntax Description	instance	Displays configuration of a specific MAGP instance Range: 1-255
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4500	
	3.6.5000	Updated example
	3.6.8100	Updated example
	3.9.1000	Updated example

Example	`switch (config) # show magp`
	
	`MAGP 100:`
	`Interface vlan: 20`
	`Admin state : Enabled`
	`State : Master`
	`Virtual IP : 11.11.11.200`
	`V6 State : Master`
	`Virtual IPv6 : 2001::254`
	`Virtual MAC : aa:bb:cc:dd:ee:ff`
	
	`Associated IP Addresses:`
	`11.11.11.254`
	
	`Associated IPv6 Addresses:`
	`2001::200`	

Related Commands		
Note		

15.7.3.8 show magp interface vlan

	show magp interface vlan <id> Displays the configuration of a specific MAGP instance.	
Syntax Description	instance	MAGP instance number Range: 1-255
Default	N/A	
Configuration Mode	Any command mode	

History	3.3.4500	
	3.6.5000	Updated example
	3.6.8100	Updated example
	3.9.1000	Updated example
Example	switch (config) # show magp interface vlan 20 MAGP 100: Interface vlan: 20 Admin state : Enabled State : Master Virtual IP : 11.11.11.200 V6 State : Master Virtual IPv6 : 2001::254 Virtual MAC : aa:bb:cc:dd:ee:ff Associated IP Addresses: 11.11.11.254 Associated IPv6 Addresses: 2001::200	
Related Commands		
Notes		

15.8 DHCP Relay

Since Dynamic Host Configuration Protocol must work correctly even before DHCP clients have been configured, the DHCP server and DHCP client need to be connected to the same network.

In larger networks, this is not always practical because each network link contains one or more DHCP relay (DHCP-R) agents. These agents receive messages from DHCP clients and forward them to DHCP servers thus extending the reach of the DHCP beyond the local network.

DHCP-R is supported for IPv4 and IPv6.

DHCP-R is supported for both primary IP subnet and secondary IP subnets.

15.8.1 DHCP-R Virtual Routing and Forwarding (VRF) Auto-Helper

In some cases it is desired that DHCP-R functionality is automatically enabled to all IP interfaces in the system. For this purpose a vrf-auto-helper may be configured on a DHCP-R instance which would provide DHCP-R services automatically for each newly created interface on a VRF.

Only one instance in each VRF can have vrf-auto-helper capability. Whenever a new instance is created in a VRF, it automatically becomes a vrf-auto-helper.

It is possible to manually disable auto-helper capability for the instance. See the command “vrf-auto-helper” for more information.

15.8.2 Upstream and Downstream Interfaces

It is possible to define an interface to be downstream, upstream, or bidirectional (both downstream and upstream):

- Bidirectional interface – capable of performing downstream and upstream functionalities
- Downstream interface (default configuration) – the interface on which queries are received from clients or from other relay agents
- Upstream interface – the interface to which queries from clients and other relay agents are forwarded

15.8.3 DHCP Relay Commands

- 15.8.1 DHCP-R Virtual Routing and Forwarding (VRF) Auto-Helper
- 15.8.2 Upstream and Downstream Interfaces
- 15.8.3 DHCP Relay Commands
  - 15.8.3.1 ip dhcp relay
  - 15.8.3.2 address
  - 15.8.3.3 always-on
  - 15.8.3.4 information option
  - 15.8.3.5 vrf
  - 15.8.3.6 port
  - 15.8.3.7 use-secondary-ip
  - 15.8.3.8 vrf-auto-helper
  - 15.8.3.9 ip dhcp relay instance (config interface)
  - 15.8.3.10 clear ip dhcp relay counters
  - 15.8.3.11 ip dhcp relay information option circuit-id
  - 15.8.3.12 ipv6 dhcp relay instance
  - 15.8.3.13 ipv6 dhcp relay instance (global server)
  - 15.8.3.14 ipv6 dhcp relay instance address (destination address on interface)
  - 15.8.3.15 ipv6 dhcp relay instance interface-id option
  - 15.8.3.16 ipv6 dhcp relay instance vrf
  - 15.8.3.17 ipv6 dhcp relay instance port
  - 15.8.3.18 ipv6 dhcp relay instance interface-id option
  - 15.8.3.19 ipv6 dhcp relay instance use-secondary-ip
  - 15.8.3.20 clear ipv6 dhcp relay counters
  - 15.8.3.21 show ip dhcp relay
  - 15.8.3.22 show ip dhcp relay counters
  - 15.8.3.23 show ipv6 dhcp relay
  - 15.8.3.24 show ipv6 dhcp relay counters

15.8.3.1 ip dhcp relay

	ip dhcp relay [instance <instance-id>] no ip dhcp relay [instance <instance-id>] Enters DHCP relay instance configuration mode, and creates DHCP instance in active VRF context. The no form of the command deletes the instance and DHCP relay process corresponding to it.	
Syntax Description	instance-id	Range: 1-8
Default	N/A	
Configuration Mode	config	
History	3.6.3004	
Example	`switch (config)# ip dhcp relay instance 1` `switch (config ip dhcp relay instance 1)#`	
Related Commands		
Notes	If an instance is not specified then instance 1 is used (if nonexistent, then it is created).	

15.8.3.2 address

	address <ip-address> no address <ip-address> Configures the DHCP server IP address on a particular instance. The no form of the command deletes the DHCP server IP address.	
Syntax Description	ip-address	Valid IP unicast address of DHCP server.
Default	N/A	
Configuration Mode	config ip dhcp relay	
History	3.3.4150	
	3.6.1002	Added VRF parameter
	3.6.3004	Enhanced command for DHCP-R multi-instance
Example	`switch (config ip dhcp relay instance 1)# address 1.2.3.4`	
Related Commands	ip dhcp relay	
Notes	• Up to 16 IP addresses may be configured • To enable DHCP relay instance, at least one IP address should be configured, or always-on parameter should be turned on using the command “ip dhcp relay always-on” • The following option for running this command is also possible: ip dhcp relay instance 1 address <ip-address>. However, if an instance is not specified then instance 1 is used (if nonexistent, then it is created).	

15.8.3.3 always-on

	always-on no always-on Enables broadcast mode on a particular instance. The no form of the command disables the broadcast mode from instance.	
Syntax Description	vrf	VRF name
Default	Disabled	
Configuration Mode	config ip dhcp relay	
History	3.3.4150	
	3.6.1002	Added VRF parameter
	3.6.3004	Enhanced command for DHCP-R multi-instance
Example	`switch (config ip dhcp relay instance 1)# always-on`	
Related Commands	ip dhcp relay	
Notes	• Broadcasts DHCP requests to all interfaces with the DHCP relay agent for given VRF • In order to enable DHCP relay, at least one IP address should be configured, or always-on parameter should be turned on using this command • When DHCP servers are configured, requests are forwarded only to configured servers • The following option for running this command is also possible: ip dhcp relay instance 1 always-on. However, if an instance is not specified then instance 1 is used (if nonexistent, then it is created).	

15.8.3.4 information option

	information option no information option Enables DHCP relay agents to insert option 82 on the packets of a particular instance. The no form of the command removes option 82 from the packets.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config ip dhcp relay	
History	3.3.4150	
	3.6.3004	Enhanced command for DHCP-R multi-instance
Example	`switch (config ip dhcp relay instance 1)# information option`	
Related Commands	ip dhcp relay	
Notes	The following option for running this command is also possible: ip dhcp relay instance 1 information option. However, if an instance is not specified then instance 1 is used (if nonexistent, then it is created).	

15.8.3.5 vrf

	vrf <vrf-name> no vrf <vrf-name> Configures mention instance in the given VRF. The no form of the command moves the instance back to default VRF.
Syntax Description	N/A
Default	N/A
Configuration Mode	config ip dhcp relay
History	3.6.3004
Example	`switch (config ip dhcp relay instance 1)# vrf 2`
Related Commands	
Notes	• If no VRF is specified, then the DHCP-R instance is created in the active VRF • If the VRF is changed, then the configuration of the DHCP-R instance is automatically deleted • The following option for running this command is also possible: ip dhcp relay instance 1 vrf <vrf-name>. However, if an instance is not specified then instance 1 is used (if nonexistent, then it is created).

15.8.3.6 port

	port <udp-port> no port <udp-port> Changes the UDP port for the given instance. The no form of the command sets the UDP port to default value.	
Syntax Description	udp-port	UDP port Range: 1-65534
Default	67	
Configuration Mode	config ip dhcp relay	
History	3.6.3004	
Example	`switch (config ip dhcp relay instance 1)# port 65534`	
Related Commands		
Notes	• The system allocated 2 ports: One is the server port (udp-port), and another is client port (udp-port+1) • The following option for running this command is also possible: ip dhcp relay instance 1 port <udp-port>. However, if an instance is not specified then instance 1 is used (if nonexistent, then it is created).	

15.8.3.7 use-secondary-ip

	use-secondary-ip no use-secondary-ip Enables the switch to relay a single request from the client multiple times simultaneously, with each of the IP addresses configured on the corresponding downstream interfaces as the respective gateway address (linkaddr field of IPv4 DHCP request packet). The no form of the command disables this function.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config ip dhcp relay
History	3.6.8008
Example	`switch (config ip dhcp relay instance 1)# use-secondary-ip`
Related Commands	
Notes	

15.8.3.8 vrf-auto-helper

	vrf-auto-helper no vrf-auto-helper Makes all L3 interfaces (existing/newly created) to be part of the given instance. The no form of the command resets this parameter to its default
Syntax Description	N/A
Default	N/A
Configuration Mode	config ip dhcp relay
History	3.6.3004
Example	`switch (config ip dhcp relay instance 1)# vrf-auto-helper`
Related Commands	
Notes	• Every new DHCP-R instance created in a VRF automatically becomes the VRF auto-helper if no other DHCP-R instance has been configured VRF auto-helper previously in that VRF • The following option for running this command is also possible: ip dhcp relay instance 1 vrf-auto-helper. However, if an instance is not specified then instance 1 is used (if nonexistent, then it is created).

15.8.3.9 ip dhcp relay instance (config interface)

	ip dhcp relay instance <instance-id> [downstream] [upstream] no ip dhcp relay instance <instance-id> [downstream] [upstream] Enables the given interface to listen for DHCP packets coming from specified instance (i.e. binds interface to that instance). The no form of the command removes the interface mapping from that instance.	
Syntax Description	instance-id	DHCP instance ID Range: 1-8

	downstream	The interface on which queries are received from clients or from other relay agents
	upstream	The interface to which queries from clients and other relay agents should be forwarded
Default	Downstream	
Configuration Mode	config interface ethernet set as router port interface config interface port-channel config interface vlan	
History	3.6.3004	
	3.6.6000	Added downstream and upstream parameters
Example	`switch (config interface ethernet 1/13)# ip dhcp relay instance 7 downstream`	
Related Commands		
Notes	• In order to enable DHCP relay, other than configuring the downstream interface, at least one IP address must be configured, or the always-on parameter must be activated using the command “ip dhcp relay always-on” • When DHCP servers are configured, requests are forwarded only to configured servers • At most, 64 interfaces can be configured on each instance • Only an existent DHCP-R may be specified • Each interface is either upstream, downstream, or bidirectional • If only downstream interfaces are defined, all interfaces in VRF are assumed to be upstream interfaces	

15.8.3.10 clear ip dhcp relay counters

	clear ip dhcp relay counters [vrf {<vrf-name> | all} | instance <instance-id>] Clears all DHCP relay counters (all interfaces) in a given VRF or instance.	
Syntax Description	vrf-name	VRF name or “all” for all VRFs
	instance-id	DHCP instance ID Range: 1-8
Default	N/A	
Configuration Mode	config	
History	3.3.4150	
	3.6.1002	Added VRF parameter
	3.6.3004	Enhanced command for DHCP-R multi-instance
	3.6.5000	Added “all” parameter
Example	`switch (config)# clear ip dhcp relay counters`	
Related Commands		
Notes	• If no DHCP-R instance is specified, then the counters of all DHCP-R instances are cleared • If a VRF is specified, then the counters of all instances on that VRF are cleared • The command “clear counters all” may also be used to clear all DHCP-R counters	

15.8.3.11 ip dhcp relay information option circuit-id

	ip dhcp relay information option circuit-id <label> no ip dhcp relay information option circuit-id Specifies the content of the circuit ID sub-option attached to the client DHCP packet when it is forwarded a DHCP server. The no form of the command removes the label assigned.	
Syntax Description	label	Specifies the label attached to packets. The string may be up to 15 characters.
Default	The label is taken from the IP interface name (e.g. “vlan1”)	
Configuration Mode	config interface vlan config interface ethernet set as router port interface config interface port-channel set as router port interface	
History	3.3.4150	
	3.6.1002	Added VRF parameter
Example	`switch (config interface vlan 10)# ip dhcp relay information options circuit-id my-label`	
Related Commands		
Notes	The circuit ID sub-option is an IP interface attribute which is shared across all DHCP-R instances.	

15.8.3.12 ipv6 dhcp relay instance

	ipv6 dhcp relay instance <instance-id> [vrf-auto-helper] [downstream] [upstream] no ipv6 dhcp relay instance <instance-id> [vrf-auto-helper] Enables DHCP relay instance configuration mode, and creates DHCP instance in active VRF context. The no form of the command deletes the DHCP relay instance.	
Syntax Description	instance-id	DHCP instance ID Range: 1-8
	vrf-auto-helper	Instance becomes VTF auto helper
	downstream	The interface on which queries are received from clients or from other relay agents
	upstream	The interface to which queries from clients and other relay agents should be forwarded
Default	Disabled	
Configuration Mode	config interface ethernet config interface port-channel config interface vlan	
History	3.6.4070	
	3.6.6000	Added downstream and upstream parameters
Example	`switch (config interface ethernet 1/1) # ipv6 dhcp relay instance 1 downstream`	
Related Commands		

Notes	• An instance without an assigned addresses is sent to All_DHCP_servers address
	• Each interface is either upstream, downstream, or bidirectional
	• At most, 64 interfaces can be configured on each instance
	• If only downstream interfaces are defined, all interfaces in VRF are assumed to be upstream interfaces
	• An instance must meet two conditions to become active:
	  • A server address or an upstream interface
	  • A downstream interface

15.8.3.13 ipv6 dhcp relay instance (global server)

	ipv6 dhcp relay instance <instance-id> address <ipv6-address or list of addresses> no ipv6 dhcp relay instance <instance-id> address <ipv6-address or list of addresses> Configure the server address on a particular instance. The no form of the command will delete the server address from instance.	
Syntax Description	instance-id	DHCP instance ID Range: 1-8
	ipv6-address	Valid global unicast IPv6 server address Up to 16 addresses can be assigned per instance
Default	N/A	
Configuration Mode	config	
History	3.6.4070	
Example	`switch (config)# ipv6 dhcp relay instance 1 address 2001::1`	
Related Commands		
Notes	An instance without an assigned addresses will send to All_DHCP_servers address	

15.8.3.14 ipv6 dhcp relay instance address (destination address on interface)

	ipv6 dhcp relay instance <instance-id> address <link-local-address> no ipv6 dhcp relay instance <instance-id> address <link-local-address> Configures the destination address on a particular instance on a specific upstream interface. Only link local address is supported. The no form of the command deletes the destination address on a specific upstream interface from a particular instance.	
Syntax Description	instance-id	DHCP instance ID Range: 1-8
	ipv6-address	Destination unicast or multicast address Only link local address in supported
Default	N/A	
Configuration Mode	config interface ethernet config interface port-channel config interface vlan	
History	3.6.4070	

Example	`switch (config interface ethernet 1/13)# ipv6 dhcp relay instance 1 address fe80::1`
Related Commands	
Notes	Up to 16 addresses can be assigned per instance

15.8.3.15 ipv6 dhcp relay instance interface-id option

	ipv6 dhcp relay instance <instance-id> interface-id option no ipv6 dhcp relay instance <instance-id> interface-id option Enables the instance to insert interface ID option. The no form of the command disables this option.	
Syntax Description	instance-id	DHCP instance ID Range: 1-8
Default	Default interface-id is an interface name (e.g. vlan1, eth1/1)	
Configuration Mode	config	
History	3.6.4070	
Example	`switch (config)# ipv6 dhcp relay instance 1 interface-id option`	
Related Commands		
Notes		

15.8.3.16 ipv6 dhcp relay instance vrf

	ipv6 dhcp relay instance <instance-id> vrf <vrf-name> no ipv6 dhcp relay instance <instance-id> vrf <vrf-name> Configures instance in the given VRF. The no form of the command will reset the instance back to default VRF.	
Syntax Description	instance-id	DHCP instance ID Range: 1-8
	vrf-name	Name of VRF
Default	Default VRF	
Configuration Mode	config	
History	3.6.4070	
Example	`switch (config)# ipv6 dhcp relay 1 vrf test`	
Related Commands		
Notes	When an instance is moved from one VRF to another - it loses all its current configuration.	

15.8.3.17 ipv6 dhcp relay instance port

	ipv6 dhcp relay instance <instance-id> port <udp-port> no ipv6 dhcp relay instance <instance-id> port <udp-port> Modifies the UDP port for the given instance. The no form of the command will set the UDP port to default value.

Syntax Description	instance-id	DHCP instance ID Range: 1-8
	port	UDP Port ID Range: 1-65534
Default	UDP port 547	
Configuration Mode	config	
History	3.6.4070	
Example	`switch (config)# ipv6 dhcp relay 1 port 555`	
Related Commands		
Notes		

15.8.3.18 ipv6 dhcp relay instance interface-id option

	ipv6 dhcp relay instance <instance-id> interface-id option [user-defined-id] Specifies the content of the interface-id option that will be sent by the relay agent.	
Syntax Description	instance-id	DHCP instance ID Range: 1-8
	user-defined-id	Interface ID option content Length: 1-15 (char) Default: interface name
Default	N/A	
Configuration Mode	config	
History	3.6.4070	
Example	`switch (config)# ipv6 dhcp relay instance <instance-id> interface-id option eth1/1`	
Related Commands		
Notes		

15.8.3.19 ipv6 dhcp relay instance use-secondary-ip

	ipv6 dhcp relay instance use-secondary-ip no ipv6 dhcp relay instance use-secondary-ip Enables the switch to relay a single request from the client multiple times simultaneously, with each of the IP addresses configured on the corresponding downstream interfaces as the respective gateway address (giaddr field of IPv6 DHCP request packet). The no form of the command disables this function.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.6.8008
Example	`switch (config ipv6 dhcp relay instance 1)# use-secondary-ip`

Related Commands	
Notes	

15.8.3.20 clear ipv6 dhcp relay counters

	clear ipv6 dhcp relay counters [vrf {<vrf-name> | all} | instance <instance-id>] Clears DHCP relay counters for specific instance or all instances in given VRF or all instances in the system.	
Syntax Description	vrf-name	VRF name or “all” for all VRFs
	instance-id	DHCP instance ID Range: 1-8
Default	N/A	
Configuration Mode	config	
History	3.6.4070	
	3.6.5000	Added “all” parameter
Example	`switch (config)# clear ipv6 dhcp relay counters vrf all`	
Related Commands		
Notes		

15.8.3.21 show ip dhcp relay

	show ip dhcp relay [instance <instance-id>] Displays general DHCP configuration.	
Syntax Description	instance-id	If instance ID is specified, then a particular instance configuration is displayed
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4150	
	3.6.1002	Added VRF and all parameters
	3.6.3004	Updated example and parameters
	3.6.6000	Updated example
	3.6.8008	Updated example

Example	switch (config)# show ip dhcp relay
	
	Instance ID 1:
	  VRF Name: default
	
	  DHCP Servers:
	  1.1.1.1
	
	  DHCP relay agent options:
	  always-on : Disabled
	  Information Option: Disabled
	  UDP port : 67
	  Auto-helper : Disabled
	
	  ---------- 	
	  Interface Label Mode
	  ---------- 	
	  eth1/5 N/A downstream

Related Commands	
Notes	• If no DHCP-R instance is given, then all DHCP-R instances are displayed • Only configured interfaces are displayed • Once vrf-auto-helper is enabled, no interface is displayed

15.8.3.22 show ip dhcp relay counters

	show ip dhcp relay counters [instance <instance-id> | vrf <vrf-name>] Displays the DHCP relay counters.	
Syntax Description	instance-id	Displays the DHCP relay counters for a given instance
	vrf	Displays the DHCP relay counters in a given VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4150	
	3.6.1002	Added VRF and all parameters
	3.6.5000	Updated example
	3.6.8008	Updated example
Example		

```
switch (config) # show ip dhcp relay counters

Instance 1:
  VRF Name: vrf-default

  DHCP Counter flags:
    SPR : Server Packets Received
    SPE : Server Packets Error
    SPRE: Server Packet Relayed
    CPR : Client Packets Received
    RP  : Relay Packets
    RE  : Relay Errors

  ----------------------------------
  Req/Resp    Received    Forwarded
  ----------------------------------
  All Req     0           0
  All Res     0           0

  -----------------------------------------------------
  If         SPRE       SPE        SPR        CPR
  -----------------------------------------------------
  eth1/5     0          0          0          0

  Packets Relayed to Server:
    ------------------------------------
    Server      RP          RE
    ------------------------------------
    1.1.1.1     0           0
```

Related Commands	
Notes	

15.8.3.23 show ipv6 dhcp relay

	show ipv6 dhcp relay [instance <instance-id>] Displays general DHCP configuration on all instances. If instance ID is defined then specific instance configuration is displayed.	
Syntax Description	instance-id	DHCP instance ID Range: 1-8
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4070	First release
	3.6.5000	Updated example
	3.6.6000	Updated example
	3.6.8008	Updated example

Example	switch (config)# show ipv6 dhcp relay
	
	Instance ID 1:
	  VRF Name: default
	
	  DHCP Servers:
	  2001:db8:701f::8f9
	
	  DHCP relay agent options:
	  All_DHCP_Servers : Disabled
	  Interface-id Option: Disabled
	  UDP port : 547
	  Auto-helper : Disabled
	  Status : Down
	
	  ---------- 	
	  Interface Label Mode
	  ---------- 	
	  eth1/5 N/A downstream

Related Commands	
Notes	• If no DHCP-R instance is given, then all DHCP-R instances are displayed • Only configured interfaces are displayed • Once vrf-auto-helper is enabled, no interface is displayed

15.8.3.24 show ipv6 dhcp relay counters

	show ipv6 dhcp relay counters [instance <instance-id> | vrf <vrf-name>] Displays the DHCPv6 relay counters.	
Syntax Description	instance-id	Displays the DHCPv6 relay counters for a given instance
	vrf	Displays the DHCPv6 relay counters in a given VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.4150	
	3.6.8008	Updated example
Example		

```
switch (config) # show ipv6 dhcp relay counters

Instance 1:
  VRF Name: vrf-default

  DHCP Counter flags:
    SPR : Server Packets Received
    SPE : Server Packets Error
    SPRE: Server Packet Relayed
    CPR : Client Packets Received
    RP  : Relay Packets
    RE  : Relay Errors

  ----------------------------------
  Req/Resp    Received    Forwarded
  ----------------------------------
  All Req     0           0
  All Res     0           0

  -----------------------------------------------------
  If         SPRE       SPE        SPR        CPR
  -----------------------------------------------------
  eth1/5     0          0          0          0

  Packets Relayed to Server:
    ---------------------------------------------------------------------
    Server                                    RP          RE
    ---------------------------------------------------------------------
    2001:db8:701f::8f9                        0           0
```

Related Commands	
Notes	

16 RDMA Over Converged Ethernet (RoCE)

16.1 RoCE Overview

RDMA over Converged Ethernet (RoCE) is a network protocol that leverages Remote Direct Memory Access (RDMA) capabilities to accelerate communications between applications hosted on clusters of servers and storage arrays. RoCE incorporates the IBTA RDMA semantics to allow devices to perform direct memory-to-memory transfers at the application level without involving the host CPU. Both the transport processing and the memory translation and placement are performed by the hardware which enables lower latency, higher throughput, and better performance compared to software-based protocols.

RoCE traffic can take advantage of IP/Ethernet L3/L2 Quality of Service (QoS). Given some of the most prevalent use cases for RDMA technology (e.g. low latency, high bandwidth), the use of QoS becomes particularly relevant in a converged environment where RoCE traffic shares the underlying network with other TCP/UDP packets. In this regard, RoCE traffic is no different than other IP flows: QoS is achieved through proper configuration of relevant mechanisms in the fabric.

RoCE Packet Structure

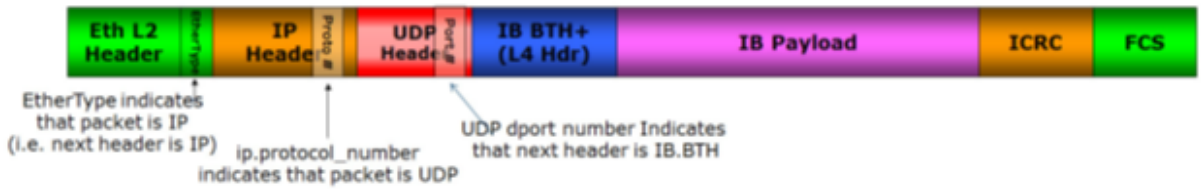

Configuration of IP/Ethernet L3/L2 QoS is determined by the RoCE application using the The SL component in the Address Vector.

RoCE Congestion Management

RoCE Congestion Management (RCM) relies on the mechanism defined in RFC3168 in the ECN protocol for the signaling of congestion. While ECN marks packets that arrive to their destination, the congestion notification is sent back to the source using a CNP packet, which limits the rate of the packet injection for the relevant QP.

16.1.1 Definitions/Abbreviation

Definitions/ Abbreviation	Description
RDMA	Remote Direct Memory Access
RoCE	RDMA over Converged Ethernet
Lossless Network	As with RoCE, the underlying networks for RoCEv2 should be configured as lossless. In this context, lossless does not mean that packets are absolutely never lost.
RCM	RoCE Congestion Management
ECN	Explicit Congestion Notification

Definitions/ Abbreviation	Description
CNP	Congestion Notification Packet
PFC	Priority Flow Control

16.2 Configuring RoCE

Configuring simplified RoCE in NVIDIA Onyx allows the user to select the RoCE configuration that best suits their use-case. To configure the simplified RoCE setting, configure the default mode of RoCE based on the NVIDIA recommended definitions or the advanced mode for specific DCN and use cases. There are three modes in which RoCE can be configured: lossless, semi-lossless, and lossy.

RoCE Configuration Modes

Options	Functionality
Lossless	This is the most optimal and automated option and is the default mode for the command, but requires a lossless network (PFC). In addition to the PFC control that exists in semi-lossless, it includes that following features: • Adds traffic pool for lossless and map switch priority (3) • Enable PFC on priority RoCE (3) on all ports.
Semi-lossless	Requires a one-way PFC between the host and the ToR (the fabric will remain lossy). In addition to the elements common to all options, it includes the following: • Micro-burst absorption (pause rx compliant, no pause propagation).
Lossy	No PFC, but has the factors common to all modes.

The following configuration is used in each of the predefined modes:

RoCE Parameters

Parameters	Lossy	Semi-lossless	Lossless
Port trust mode L3	✓	✓	✓
Port sw-prio-TC mapping • **sw-prio 3**—TC 3 (RoCE) • **sw-prio 6**—TC 6 (CNP) • other sw-prio—TC 0	✓	✓	✓
Port ETS • **TC 6 (CNP)**—strict • **TC 3 (RoCE)**—WWR 50% • **TC 0 (other traffic)**—WWR 50%	✓	✓	✓
Port ECN absolute threshold 150-1500 **TC 3 (RoCE)**	✓	✓	✓
LLDP + Application TLV (RoCE) (UDP, Protocol: 4791, Priority 3)	✓	✓	✓
Enable PFC on sw-prio 3 (RoCE)		✓	✓

Parameters	Lossy	Semi-lossless	Lossless
Prio 3 to roce lossless traffic pool			✓

- The RoCE command defines the switch default values for several parameters defined in details in the RoCE Parameters table, above. Changes made by the user for RoCE-related parameters will not be changed by the RoCE command when executed.
- Changing buffer configuration mode to "advanced buffer management" after configuring RoCE returns the buffer configuration to its default configuration.

16.3 RoCE Commands

- RoCE Commands

16.4 Further Information

For more information about this feature and its potential applications, please refer to the following community posts:

- How To Enable, Verify and Troubleshoot RDMA
- RDMA/RoCE Solutions
- RoCE v2 Considerations
- How To add a Timestamp in RoCE
- Understanding RoCEv2 Congestion Management
- MTU Considerations for RoCE based Applications
- Recommended Network Configuration Examples for RoCE Deployment
- How To Configure RoCEv2 for ConnectX-3 Pro Using SwitchX Switches
- Understanding QoS Configuration for RoCE
- How To Configure RoCE Over a Lossless Fabric (PFC+ECN) End-to-End Using ConnectX-4 and Spectrum (Trust L2)
- How To Run RoCE Over L2 Enabled With PFC
- Lossless RoCE Configuration for Onyx Switches in DSCP-Based QoS Mode
- How To Configure RoCE Over a Lossy Fabric (ECN) End-to-End Using ConnectX-4 and Spectrum (Trust L3)
- How To Configure RoCE With ECN End-to-End Using ConnectX-4 and Spectrum (Trust L2)
- RoCE Configuration for Onyx Switches in PCP-Based QoS Mode (Advanced Mode)
- How To Configure Resilient RoCE End-to-End Using ConnectX-4 and Spectrum (No QoS)
- Lossless RoCE Configuration for Onyx Switches in PCP-Based QoS Mode
- How To Configure Spectrum Switch for Lossless RoCE

- How To Configure Spectrum Switch for Resilient RoCE
- RoCE Configuration for Onyx Switches in DSCP-Based QoS Mode
- Lossless RoCE Configuration for MLNX-OS Switches in DSCP-Based QoS Mode (Advanced Mode)

16.5 RoCE Commands

- 16.5.1 roce
- 16.5.2 show roce
- 16.5.3 show interfaces ethernet 1/1 counters roce
- 16.5.4 clear roce interface ethernet 1/1

16.5.1 roce

	roce [< lossy | semi-lossless | lossless >] [no] roce Configures the switch to RoCE mode. The no form of the command disables RoCE mode.	
Syntax Description	Lossless	Full PFC support (this is the default when no parameter is chosen).
	Semi-lossless	Micro-burst absorption (pause rx compliant, no pause propagation).
	Lossy	Congestion control based on ECN marking only. No PFC support.
Default	N/A	
Configuration Mode	config	
History	3.8.2000	

Example	`switch (config) # roce <mode>`
	
	`switch (config) # no roce`
	`switch (config) # show roce`
	
	`RoCE mode: N/A`
	
	`switch (config) #`	

Related Commands	`show roce` `show interfaces ethernet 1/1 counters roce`	
Notes	• Configuring RoCE without specifying a mode will configure RoCE with lossless mode. • Changing RoCE mode may cause interfaces toggling and, consequently, a momentary loss of data.	

16.5.2 show roce

	show roce Displays RoCE mode information.
Syntax Description	N/A

Default	N/A	
Configuration Mode	config	
History	3.8.2000	
	3.8.3000	Updated example
Example		

```
switch (config) # show roce

RoCE mode      : lossless
LLDP           : disabled
Port trust mode: L3

Application TLV:
  Selector: udp
  Protocol: 4791
  Priority: 3

Port congestion-control:
  Mode: ecn, absolute
  Min : 150
  Max : 1500

PFC              : enabled
switch-priority 3: enabled

RoCE used TCs:
  ------------------------------------------
  Switch-Priority   TC     Application   ETS
  ------------------------------------------
  3                 3      RoCE          WRR 50%
  6                 6      CNP           Strict

RoCE buffer pools:

-------------------------------------------------------------------------------
-----------
  Traffic                  Type      Memory   Switch        Memory actual   Usage
Max Usage
  Pool                               [%]      Priorities

-------------------------------------------------------------------------------
-----------
  lossy-default            lossy     auto     0, 1, 2, 5,   14.4M           0
0
                                              6, 7
  roce-reserved            lossless  auto     3, 4          14.4M           0
0

Exception list:
Switch priority 4 is mapped to RoCE traffic pool
LLDP is not enabled.
Interface ethernet 1/8 PFC is not enabled.

Json output:
[
    {
        "LLDP": "disabled",
        "Port trust mode": "L3",
        "RoCE mode": "lossless"
    },
    {
        "Application TLV": [
            {
                "Priority": "3",
                "Protocol": "4791",
                "Selector": "udp"
            }
        ]
    },
    {
        "Port congestion-control": [
            {
                "Max": "1500",
                "Mode": "ecn, absolute",
                "Min": "150"
            }
        ]
```

```
    },
    {
        "PFC": "enabled",
        "switch-priority 3": "enabled"
    },
    {
        "RoCE used TCs": [
            {
                "3": [
                    {
                        "Application": "RoCE",
                        "TC": "3",
                        "ETS": "WRR 50%"
                    }
                ],
                "6": [
                    {
                        "Application": "CNP",
                        "TC": "6",
                        "ETS": "Strict"
                    }
                ]
            }
        ]
    },
    {
        "RoCE buffer pools": [
            {
                "roce-reserved": [
                    {
                        "Type": "lossless",
                        "Switch Priorities": "3, 4",
                        "Max Usage": "0",
                        "Usage": "0",
                        "Memory actual": "14.4M",
                        "Memory [%]": "auto"
                    }
                ],
                "lossy-default": [
                    {
                        "Type": "lossy",
                        "Switch Priorities": "0, 1, 2, 5, 6, 7",
                        "Max Usage": "0",
                        "Usage": "0",
                        "Memory actual": "14.4M",
                        "Memory [%]": "auto"
                    }
                ]
            }
        ]
    },
    {
        "Exception list": [
            {
                "Lines": [
                    "Switch priority 4 is mapped to RoCE traffic pool",
                    "LLDP is not enabled.",
                    "Interface ethernet 1/8 PFC is not enabled."
                ]
            }
        ]
    }
```


Related Commands	show roce show interfaces ethernet 1/1 counters roce
Notes	Interface-related properties (such as ETS, QoS, TC mapping) represent expected values for RoCE. For the state of a specific interface, please use relevant interface show command.

16.5.3 show interfaces ethernet 1/1 counters roce

	show interfaces ethernet 1/1 counters roce Display specific interfaces counters relevant to RoCE. See example below.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.8.2000
Example	switch (config) # show interfaces ethernet 1/1 counters roce Rx: 0 RoCE PG packets 0 RoCE PG bytes 0 RoCE no buffer discard 0 CNP PG packets 0 CNP PG bytes 0 CNP no buffer discard 0 RoCE PFC pause packets 0 RoCE PFC pause duration 0 RoCE buffer usage (bytes) 0 RoCE buffer max usage (bytes) 0 CNP buffer usage (bytes) 0 CNP buffer max usage (bytes) 0 RoCE PG usage (bytes) 0 RoCE PG max usage (bytes) 0 CNP PG usage (bytes) 0 CNP PG max usage (bytes) Tx: 0 ECN marked packets 0 RoCE TC packets 0 RoCE TC bytes 0 RoCE unicast no buffer discard 0 CNP TC packets 0 CNP TC bytes 0 CNP unicast no buffer discard 0 RoCE PFC pause packets 0 RoCE PFC pause duration 0 RoCE buffer usage (bytes) 0 RoCE buffer max usage (bytes) 0 CNP buffer usage (bytes) 0 CNP buffer max usage (bytes) 0 RoCE TC usage (bytes) 0 RoCE TC max usage (bytes) 0 CNP TC usage (bytes) 0 CNP TC max usage (bytes)
Related Commands	roce show roce
Notes	

16.5.4 clear roce interface ethernet 1/1

	clear roce interface ethernet 1/1 Clears all the counters including the max-usage counters.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.8.2000

Example	`switch (config) # clear roce interface ethernet 1/1`
Related Commands	`show interfaces ethernet 1/1 counters roce` clear counters clear buffers interface ethernet 1/1 max-usage
Notes	

17 Multicast (IGMP and PIM)

Protocol independent multicast (PIM) is a collection of protocols that deal with efficient delivery of IP multicast (MC) data. Those protocols are published in the series of RFCs and define different ways and aspects of multicast data distribution. PIM protocol family includes Internet Group Management protocol (IGMP), IGMP Snooping, Bootstrap router (BSR) protocol, and PIM variations: Sparse mode (PIM-SM), Source-Specific mode (PIM-SSM), Dense mode (PIM-DM) and Bidirectional mode (PIM-BIDIR). PIM-DM in not supported in Onyx.

PIM builds and maintains multicast routing tables based on the unicast routing information provided by unicast routing tables that can be maintained statically or dynamically by IP routing protocols like OSPF and BGP.

17.1 Basic PIM-SM

PIM relies on the underlying topology gathering protocols that collect unicast routing information and build multicast routing information base (MRIB). The primary role of MRIB is to determine the next hop for PIM messages. MC data flows along with the reverse path of the PIM control.

MC tree construction contains three phases:

1. Construction of a shared distribution tree. This tree is built around a special router called the rendezvous point (RP).
2. Establishing a native forwarding path from MC sources to the RP.
3. Building an optimized MC distribution tree directly from each MC source to all MC targets.

The first stage of the multicast tree establishment starts when the MC receiver expresses desire to start receiving MC data. It can happen as a result of using one of the L3 protocols like MLD or IGMP, or by static configuration. When such request is received by the last hop router (a designated router) this router starts to build a distribution path from the RP. It starts to send periodic “Join” messages to the nearest PIM neighbor router towards the RP. The next router continues to do the same. Eventually the process converges when Join messages reach RP or a router that has already created that distribution tree. Usually that tree is called a shared tree because it is created for any source for specific MC group G and is noted as (*,G).

At that stage, MC senders can start sending MC data. The DR next to the MC source extracts the packets from the data flow and tunnels them to the RP. The RP decapsulates the packets and distributes them to all MC receivers along with the share tree.

On the second stage the RP switches from tunneling of multicast packets from MC sources to forwarding native traffic. When the RP identifies that a new MC source started to send packets, it initiates an establishment of a native forwarding path from the DR of that source to itself. For this purpose it starts to send Join messages towards MC source to nearest neighbor to that source according the MRIB. This is a source specific Join and is noted as (S,G). When data path is established up to the DR, the DR switches from tunneling MC packets to their native forwarding, so the RP does not need to decapsulate MC packets anymore, but still continue to distribute the packets along with shared tree.

On the third phase multicast receivers will try to switch from shared tree to source specific tree by creating a direct distribution path from a multicast source. When last hop router of the multicast receiver identifies multicast traffic coming from any multicast source it will start to send Join messages towards the source with purpose to create a direct source specific path to that source. Once such path will be established and Designated router that is attached to the source L2 network

will start to distribute the multicast traffic directly bypassing shared tree, the last hop router will detach its receivers from shared tree for that data and will switch to the shortest path tree distribution.

17.2 Source-Specific Multicast (SSM)

Source-Specific Multicast (SSM) is a method of delivering multicast packets in which the only packets that are delivered to a receiver are those originating from a specific source address requested by the receiver. By so limiting the source, SSM reduces demands on the network and improves security.

SSM requires that the receiver specify the source address and explicitly excludes the use of the (*,G) join for all multicast groups in RFC 3376, which is possible only in IPv4's IGMPv3 and IPv6's MLDv2.

Source-specific multicast is best understood in contrast to any-source multicast (ASM). In the ASM service model a receiver expresses interest in traffic to a multicast address. The multicast network must discover all multicast sources sending to that address, and route data from all sources to all interested receivers.

This behavior is particularly well suited for groupware applications where all participants in the group want to be aware of all other participants, and the list of participants is not known in advance.

The source discovery burden on the network can become significant when the number of sources is large.

In the SSM service model, in addition to the receiver expressing interest in traffic to a multicast address, the receiver expresses interest in receiving traffic from only one specific source sending to that multicast address. This relieves the network of discovering many multicast sources and reduces the amount of multicast routing information that the network must maintain.

SSM requires support in last-hop routers and in the receiver's operating system. SSM support is not required in other network components, including routers and even the sending host. Interest in multicast traffic from a specific source is conveyed from hosts to routers using IGMPv3 as specified in RFC 4607.

By default SSM destination addresses defined in the ranges 232.0.0.0/8 for IPv4 or FF3x::/96 for IPv6. This range may be configured by user.

Source-specific multicast delivery semantics are provided for a datagram sent to an SSM address. That is, a datagram with source IP address S and SSM destination address G is delivered to each upper-layer “socket” that has specifically requested the reception of datagrams sent to address G by source S, and only to those sockets.

17.3 Bidirectional PIM

Bidirectional PIM (PIM-BIDIR) is a variant of PIM-SM that builds bidirectional distribution trees that connect multicast senders and receivers. It differs from PIM-SM by eliminating a need to tunnel multicast packets to RP and to keep a state for each (S,G) pair. It also eliminates a need in data driven protocol events. PIM-BIDIR achieves it by defining a new role, Designated Forwarder (DF), and by defining new forwarding rules and keeping all other PIM-SM mechanisms intact.

DF is a PIM enabled router that is the closest router to RP among all PIM routers residing on specific L2 network. It is dynamically elected by all PIM routers on that network. DF is required on each L2

multicast capable network for each RP. DF serves all multicast groups that share the same RP and has following duties:

- It is an only router that is responsible to receive and forward upstream multicast packets on that L2 segment
- It is a router that should collect all Join requests from the routers on that L2 segment
- It is an only router that will distribute downstream multicast packets on that segment.

Once Designated forwarders are elected and forwarding rules are established, PIM routers can start to issue (*,G) Join messages and build shared distribution trees. When shared tree is created, multicast sources can start to exchange data with receivers and it doesn't require any additional maintenance of the multicast states.

Compared to PIM-SM, in bidirectional PIM:

- Each router will keep only (*,G) state and not (*,G) and (S,G) like in PIM-SM
- Multicast traffic from the beginning is forwarded naturally - no need to tunnel data to RP
- Resulting multicast tree is not shortest path optimal and converges around selected Rendezvous point, but is shared among all participants in that multicast group

In BIDIR-PIM, the packet forwarding rules have been improved over PIM-SM, allowing traffic to be passed up the shared tree toward the RP. To avoid multicast packet looping, bidir-PIM introduces a new mechanism called designated forwarder (DF) election, which establishes a loop-free SPT rooted at the RP.

17.4 PIM Load-Sharing

PIM load-sharing improves network efficiency in IP multicast applications especially in cases when we have multiple equal-cost paths to the same destination. There two methods which enhance IP multicast bandwidth capacity consumption: rendezvous point load sharing and next-hop load sharing.

> Routers should be connected via router port interfaces and not VLAN interfaces. Connecting two routers via VLAN interface with PIM load-sharing causes loops in the network.

17.4.1 Rendezvous Point Load-Sharing

IP multicast routing is facilitated by use of rendezvous points (RPs) which are anchors in IP multicast distribution trees, and, in case of PIM-BIDIR, are central points that perform IP multicast packet forwarding. Therefore, they can get heavily loaded.

When multiple RPs serve the same multicast IP addresses and are located at an equal distance from a traffic source or receiver, data streams can be shared between those RPs. This enhances switching performance, improves network bandwidth consumption and increases reliability. Data packets based on the packet flow parameters are equally shared between all RPs located at an equal-distance.

17.4.2 Next Hop Load-Sharing

Another way to improve network capacity consumption and increase the amount of IP multicast data carried by the network, is to utilize multiple equal-cost paths from RPs to IP multicast receivers. A network usually selects a single path to carry specific multicast group data packets from a source to a specific multicast destination. But when enabling next hop load-sharing, multiple paths between RP and multicast group receivers may be utilized, and based on traffic flow parameters, the data stream may be split to multiple flows that go through several equal-cost paths to the same destination.

17.5 Bootstrap Router

For correct operation each PIM router requires a capability to map a multicast group that it needs to serve to a Rendezvous point for that group. This mapping can be done manually or the mapping can be distributed dynamically in the network. BSR protocol serves for this purpose.

This protocol introduces new role in the multicast network – Bootstrap router. That router is responsible to flood multicast group to RP mapping through the multicast routing domain. Bootstrap router is elected dynamically among bootstrap router candidates (C-BSR) and once elected will collect from Rendezvous point candidate (C-RP) mapping information and distribute it in the domain.

Bootstrap activity contains 4 steps. First each C-BSR configured in the network originates floods into the network bootstrap messages that express the router desire to become BSR and also its BSR priority. Any C-BSR that receives that information and has lower priority will suspend itself, so eventually only one router will send BSR messages and become BSR.

When BSR is elected all RP candidates start to advertise to BSR a list of groups that this RP can serve. On the next step, after BSR learns the group mapping proposals, it forms a final group to RP mapping in the domain and starts to distribute it among PIM routers in the multicast routing domain. When PIM router receives BSR message with the group to RP mapping, it installs that mapping in the router local cache and uses that information to create multicast distribution trees.

17.6 Configuring Multicast

Precondition steps:

1. Enable IP routing functionality. Run:

```
switch (config)# ip routing
```

2. Enable the desired VLAN. Run:

```
switch (config)# vlan 10
```

3. Add this VLAN to the desired interface. Run:

```
switch (config)# interface ethernet 1/1
switch (config interface ethernet 1/1)# switchport access vlan 10
```

4. Create a VLAN interface. Run:

```
switch (config)# interface vlan 10
```

5. Apply IP address to the VLAN interface. Run:

```
switch (config interface vlan 10)# ip address 10.10.10.10 /24
```

6. Enable the interface. Run:

```
switch (config interface vlan 10)# no shutdown
```

17.6.1 Configuring IGMP

IGMP is enabled when IP multicast is enabled and static multicast or PIM is enabled on the interface.

17.6.2 Verifying IGMP

1. Display a brief IGMP interface status. Run:

```
switch (config)# show ip igmp interface brief

VRF "default":

  ---------------------------------------------------------------------------
  Interface     IP Address      IGMP Querier     Membership Count     Version
  ---------------------------------------------------------------------------
  Vlan10        10.10.10.1      10.10.10.1       1                    v2
```

2. Display detailed IGMP interface status. Run:

```
switch (config)# show ip igmp interface vlan 10
Interface vlan10
        Status: protocol-down/link-down/admin-up
        VRF: "vrf-default"
        IP address: 10.10.10.1/24
        Active querier: 10.10.10.1
        Version: 2
        Next query will be sent in: 00:01:45
        Membership count: 0
        IGMP version: 2
        IGMP query interval: 125 secs
        IGMP max response time: 10 secs
        IGMP startup query interval: 31 secs
        IGMP startup query count: 2
        IGMP last member query interval: 1 secs
        IGMP last member query count: 2
        IGMP group timeout: 260 secs
        IGMP querier timeout: 0 secs
        IGMP unsolicited report interval: 10 secs
        IGMP robustness variable: 2
        IGMP interface immediate leave: Disabled
        Multicast routing status on interface: Enabled
        Multicast TTL threshold: 0

        IGMP interface statistics:
          General (sent/received):
            v2-queries: 2/0
            v2-reports: 0/0
            v2-leaves : 0/0
            v3-queries: 0/0
            v3-reports: 0/0

          Errors:
            Checksum errors                         : 0
            Packet length errors                    : 0
            Packets with Local IP as source         : 0
            Source subnet check failures            : 0
            Query from non-querier                  : 0

            Report version mismatch                 : 0
            Query version mismatch                  : 0
            Unknown IGMP message type               : 0
            Invalid v2 reports                      : 0
            Invalid v3 reports                      : 0
            Invalid leaves                          : 0
            Packets dropped due to router-alert check: 0
```

3. Display the list of IGMP groups and their status. Run:

```
switch (config)# show ip igmp groups
IGMP Connected Group Membership
Type: S - Static, D - Dynamic

------------------------------------------------------------------------------------------------------------------
------------
Group Address        Type                Interface           Uptime              Expires
Last Reporter
------------------------------------------------------------------------------------------------------------------
------------
226.0.1.0            D                   vlan10              00:00:05            N/A
10.10.10.2
226.0.1.1            D                   vlan10              00:00:04            N/A
10.10.10.2
```

17.6.3 Configuring PIM

Prerequisites:

1. If not enabled, enable IP routing. Run:

```
switch (config)# ip routing
```

2. Globally enable multicast routing. Run:

```
switch (config)# ip multicast-routing
```

To configure PIM:

1. Enable PIM. Run:

```
switch (config)# protocol pim
```

2. Enable PIM on any IP interface (router port or VLAN interface) facing an L3 multicast source or L3 multicast receiver including transit interfaces. For example, run:

```
switch (config)# interface ethernet 1/4 ip pim sparse-mode
```

> The interface's primary address is always used in PIM.

3. Configure IGMP version on any IP interface (router port or VLAN interface) facing multicast receivers. For example, run:

```
switch (config)# interface ethernet 1/4 ip igmp version {2|3}
```

If IGMP must be enabled on a VLAN interface, IP IGMP snooping must also be enabled (globally and on the relevant VLAN interface):

```
switch (config)# interface vlan 50 ip igmp version {2|3}
switch (config)# ip igmp snooping
switch (config)# vlan 50 ip igmp snooping
```

4. Configure a rendezvous point. Run:

```
switch (config)# ip pim rp-address 10.10.10.10
```

A good practice is to configure the RP on the loopback interface. Although RP may be configured on the any interface with enabled PIM sparse mode. Note that a loopback interface does not require enabling PIM sparse mode to configure RP.

The RP address must be reachable to all switches.

5. Configure a group mapping for a static RP. Run:

```
switch (config)# ip pim rp-address 192.168.0.1
```

You may also specify a “group-list <ip-address> <prefix>” parameter (ip pim rp-address 192.168.0.1 group-list 224.0.0.0/4) if you want different RPs for different groups.

17.7 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community post:

- HowTo Configure IP Multicast (PIM, IGMP) on Ethernet Switches

17.8 IGMP and PIM Commands

- IGMP and PIM Commands
- IGMP Snooping

17.9 IGMP and PIM Commands

- 17.9.1 PIM
  - 17.9.1.1 protocol pim
  - 17.9.1.2 ip pim sg-expiry-timer
  - 17.9.1.3 ip pim rp-address
  - 17.9.1.4 ip pim bsr-candidate
  - 17.9.1.5 ip pim register-source
  - 17.9.1.6 ip pim rp-candidate
  - 17.9.1.7 ip pim sparse-mode
  - 17.9.1.8 ip pim dr-priority
  - 17.9.1.9 ip pim hello-interval
  - 17.9.1.10 ip pim join-prune-interval
  - 17.9.1.11 ip pim ssm range

        - 17.9.1.12 ip pim multipath next-hop
        - 17.9.1.13 ip pim multipath rp
        - 17.9.1.14 clear ip pim counters
        - 17.9.1.15 show ip pim protocol
        - 17.9.1.16 show ip pim bsr
        - 17.9.1.17 show ip pim interface
        - 17.9.1.18 show ip pim interface brief
        - 17.9.1.19 show ip pim neighbor
        - 17.9.1.20 show ip pim rp
        - 17.9.1.21 show ip pim rp-hash
        - 17.9.1.22 show ip pim rp-candidate
        - 17.9.1.23 show ip pim ssm range
        - 17.9.1.24 show ip pim upstream joins
    - 17.9.2 PIM Bidir
        - 17.9.2.1 ip pim bidir shutdown
        - 17.9.2.2 ip pim df-robustness
        - 17.9.2.3 ip pim df-backoff-interval
        - 17.9.2.4 ip pim df-offer-interval
        - 17.9.2.5 show ip pim interface df
    - 17.9.3 Multicast
        - 17.9.3.1 ip multicast-routing
        - 17.9.3.2 ip mroute
        - 17.9.3.3 ip multicast ttl-threshold
        - 17.9.3.4 clear ip mroute
        - 17.9.3.5 show ip mroute
        - 17.9.3.6 show ip mroute summary
    - 17.9.4 IGMP
        - 17.9.4.1 ip igmp immediate-leave
        - 17.9.4.2 ip igmp last-member-query-response-time
        - 17.9.4.3 ip igmp startup-query-count
        - 17.9.4.4 ip igmp startup-query-interval
        - 17.9.4.5 ip igmp query-interval
        - 17.9.4.6 ip igmp query-max-response-time
        - 17.9.4.7 ip igmp robustness-variable
        - 17.9.4.8 ip igmp static-oif
        - 17.9.4.9 clear ip igmp groups
        - 17.9.4.10 show ip igmp groups
        - 17.9.4.11 show ip igmp interface
        - 17.9.4.12 show ip igmp interface brief

17.9.1 PIM

17.9.1.1 protocol pim

	protocol pim no protocol pim Enables protocol independent multicast (PIM). The no form of the command hides all PIM commands and deletes all PIM configurations.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.3.5006
Example	`switch (config) # protocol pim`
Related Commands	
Notes	

17.9.1.2 ip pim sg-expiry-timer

	ip pim [vrf <vrf-name>] sg-expiry-timer <seconds> no ip pim [vrf <vrf-name>] sg-expiry-timer Adjusts the SG expiry timer interval for PIM-SM SG multicast routes. The no form of the command resets the parameters to their default values	
Syntax Description	vrf	VRF name
	seconds	Range: 1-65535
Default	180 seconds	
Configuration Mode	config	
History	3.6.6102	
Example	`switch (config) # ip pim sg-expiry-timer 180`	
Related Commands		
Notes		

17.9.1.3 ip pim rp-address

	ip pim [vrf <vrf-name>] rp-address <rp-address> [group-list <ip-address> <prefix>] [override] [bidir] no ip pim [vrf <vrf-name>] rp-address <rp-address> [group-list <ip-address> <prefix>] [override] [bidir] Configures a static IP address of a rendezvous point for a multicast group range or adds new multicast range to existing RP. The no form of the command removes the rendezvous point for a multicast group range or removes all configuration of the RP.	
Syntax Description	vrf	VRF name

	rp-address	The static IP address of rendezvous point
	ip-address	IP address of the group-range (coupled with the prefix parameter)
	prefix	Network prefix (in the format of /24, or 255.255.255.0 for example) of group range
	override	Specifies that this configuration overrides dynamic configuration learned by BSR
	bidir	Optional during configuration, but appears in the configuration if in PIM Bidir mode
Default	N/A	
Configuration Mode	config	
History	3.3.5006	
	3.9.1900	Added bidir option

Example	`switch (config) # ip pim rp-address 10.10.10.10`
	
	`switch (config) # ip pim vrf default rp-address 100.100.100.100 group-list 233.3.3.3/32 bidir`	

Related Commands		
Notes		

17.9.1.4 ip pim bsr-candidate

	ip pim [vrf <vrf-name>] bsr-candidate {vlan <vlan-id> | loopback <number> | ethernet <port> | port-channel <id>} [hash-len <hash-length>] [priority <priority>] [interval <interval>] no ip pim [vrf <vrf-name>] bsr-candidate {vlan <vlan-id> | loopback <number> | ethernet <port>} [hash-len <hash-length>] [priority <priority>] [interval <interval>] Configures the switch as a candidate BSR router (C-BSR). The no form of the command removes BSR-candidate configuration or restores default parameters values.	
Syntax Description	vrf	VRF name
	vlan <vlan-id>	VLAN ID. Range: 1-4094.
	loopback <number>	Loopback interface for the BSR candidate address
	ethernet <port>	Ethernet interface for the BSR candidate address
	port-channel <id>	LAG interface for the BSR candidate address
	hash-len	Specifies the hash mask length used in BSR messages. Range: 0-32.
	priority	BSR priority rating. Larger numbers denote higher priority. Range: 0-255.
	interval	Period between the transmission of BSMs (seconds). Range: 10-536870906.
Default	The interface is not BSR candidate by default. priority—64 interval—60 hash-len—30	

Configuration Mode	config config interface ethernet (configured as a router port interface) config interface loopback config interface port-channel (configured as a router port interface) config interface vlan
History	3.3.5006
Example	`switch (config) # ip pim bsr-candidate vlan 10 priority 100`
Related Commands	ip pim sparse-mode
Notes	• A BSR is a PIM router within the PIM domain through which dynamic RP selection is implemented. The BSR selects RPs from a list of candidate RPs and exchanges bootstrap messages (BSM) with all routers in the domain. The BSR is elected from one of the C-BSRs through an exchange of BSMs. A subset of PIM routers within the domain are configured as candidate Bootstrap routers (C-BSRs). Through the exchange of Bootstrap messages (BSMs), the C-BSRs elect the BSR, which then uses BSMs to inform all domain routers of its status. • Command parameters specify the switch's BSR address, the interval between BSM transmissions, hash length used for RP calculations and the priority assigned to the switch when electing a BSR • Entering an ip pim bsr-candidate command replaces any previously configured bsr-candidate command. If the new command does not specify a priority or interval, the previously configured values persist in running-config.

17.9.1.5 ip pim register-source

	ip pim [vrf <vrf-name>] register-source <interface> no ip pim [vrf <vrf-name>] register-source <interface> Configures interface from which to use IP as source in PIM communications. The no form of the command undoes this configuration.	
Syntax Description	vrf	VRF name
	interface	Interface whose IP to use
Default	N/A	
Configuration Mode	config config interface ethernet (configured as a router port interface) config interface loopback config interface port-channel (configured as a router port interface) config interface vlan	
History	3.6.6102	
Example	`switch (config) # ip pim register-source ethernet 1/2`	
Related Commands		
Notes	This command must be set on an L3 interface with PIM sparse-mode (and not on a regular L3 interface which is not a PIM interface)	

17.9.1.6 ip pim rp-candidate

	ip pim [vrf <vrf-name>] rp-candidate {vlan <vlan-id> | loopback <number> | ethernet <slot/port>} group-list <ip-address> <prefix> [priority <priority>] [interval <interval>] [bidir] no ip pim [vrf <vrf-name>] rp-candidate {vlan <vlan-id> | loopback <number> | ethernet <slot/port>} group-list <ip-address> <prefix> [priority <priority>] [interval <interval>] [bidir] Configures the switch as a candidate rendezvous point (C-RP). The no form of the command removes the ip pim rp-candidate from running-config command for the specified multicast group.	
Syntax Description	vrf	VRF name
	ethernet <slot/port>	Ethernet interface
	port-channel <number>	LAG interface
	vlan <vlan-id>	VLAN ID Range: 1-4094
	loopback <number>	Loopback interface number
	ip-address	The group IP address
	prefix	Network prefix (for example /24, or 255.255.255.0)
	priority	RP priority rating Range: 0-255, where smaller numbers mean higher priority
	interval	RP-advertisements message transmission interval Range: 0-16383
	bidir	Optional during configuration, but appears in the configuration if in PIM Bidir mode
Default	RP priority—192 BSR message interval—60 seconds	
Configuration Mode	config config interface ethernet (configured as a router port interface) config interface loopback config interface port-channel (configured as a router port interface) config interface vlan	
History	3.3.5006	
	3.9.0300	Updated example
	3.9.1900	Added bidir option
Example	switch (config) # interface ethernet 1/13 ip pim ? bfd Configure BFD protection for PIM neighbors on the interface dr-priority Configure PIM DR priority on interface hello-interval Configure PIM hello interval on interface join-prune-interval Configure PIM join-prune interval on interface sparse-mode Configure PIM sparse mode on the interface switch (config) # interface ethernet 1/13 ip pim bfd switch (config) # ip pim vrf default rp-candidate ethernet 1/12 group-list 225.1.0.0/16 switch (config) # ip pim vrf default rp-candidate ethernet 1/12 bidir	

Related Commands	
Note	• The BSR selects a multicast group's dynamic RP set from the list of C-RPs in the PIM domain. The command specifies the interface (used to derive the RP address), C-RP advertisement interval, and priority rating. The BSR selects the RP set by comparing C-RP priority ratings. The C-RP advertisement interval specifies the period between successive C-RP advertisement message transmissions to the BSR. • Running-config supports multiple multicast groups through multiple ip pim rp-candidate statements • All commands must specify the same interface. Issuing a command with an interface that differs from existing commands removes all existing commands from running-config. • Running-config stores the interval and priority setting in a separate statement that applies to all rp-candidate statements. When a command specifies an interval that differs from the previously configured value, the new value replaces the old value and applies to all configured rp-candidate statements. • When the no commands do not specify a multicast group, all rp-candidate statements are removed from running-config. The no ip pim rp-candidate interval commands restore the interval setting to the default value of 60 seconds. • When setting a priority, all previous rp-candidates within all interfaces and groups are configured to this priority

17.9.1.7 ip pim sparse-mode

	ip pim sparse-mode no ip pim sparse-mode Sets PIM sparse mode on this interface. The no form of the command disables the sparse-mode on the interface and deletes all interfaces configuration.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)
History	3.3.5006
Example	`switch (config interface vlan 10) # ip pim sparse-mode`
Related Commands	
Notes	

17.9.1.8 ip pim dr-priority

	ip pim dr-priority <priority> no ip pim dr-priority Configures the designated router (DR) priority of PIM Hello messages. The no form of the command resets this parameter to its default.	
Syntax Description	priority	The designated router priority of the PIM Hello messages. Range is 1-4294967295.
Default	1	

Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)
History	3.3.5006
Example	`switch (config interface vlan 10) # ip pim dr-priority 5`
Related Commands	ip pim sparse-mode
Notes	The command “ip pim sparse-mode” must be run prior to using this command.

17.9.1.9 ip pim hello-interval

	ip pim hello-interval <interval> no ip pim hello-interval Configures PIM Hello interval in seconds. The no form of the command resets this parameter to its default.	
Syntax Description	interval	PIM Hello interval Range: 1-18000
Default	30 seconds	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.5006	
	3.6.4006	Updated parameter range
Example	`switch (config interface vlan 10) # ip pim hello-interval 7000`	
Related Commands	ip pim sparse-mode	
Notes	The command “ip pim sparse-mode” must be run prior to using this command	

17.9.1.10 ip pim join-prune-interval

	ip pim join-prune-interval <period> no ip pim join-prune-interval Configures the period between Join/Prune messages that the configuration mode interface originates and sends to the upstream RPF neighbor. The no form of the command resets this parameter to its default.	
Syntax Description	period	Range: 1-18000 seconds
Default	60 seconds	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.5200	
	3.6.4006	Updated parameter range
Example	`switch (config interface vlan 10) # ip pim join-prune-interval 60`	
Related Commands		
Notes		

17.9.1.11 ip pim ssm range

	ip pim [vrf <vrf-name>] ssm range {standard | group-list {<group-range>|<address> <prefix>}} no ip pim [vrf <vrf-name>] ssm range {standard | group-list {<group-range>| <address> <prefix>}} Enables one or more ranges for SSM operation. The no form of the command disables range for SSM operation.	
Syntax Description	vrf	VRF name
	standard	Sets the SSM operation to standard SSM range 232.0.0.0/8
	<group-range>	User-defined multicast range for SSM operation (e.g. 233.0.0.0/8)
	<ip-address>	Group range ip-address (e.g. 233.0.0.0/8)
	<prefix>	Group range prefix (e.g. 233.0.0.0/8)
Default	N/A	
Configuration Mode	config	
History	3.6.4006	
Example	`switch (config) # ip pim ssm range group-list 234.0.0.0/8`	
Related Commands		
Notes	Standard and group-list configurations are mutually exclusive. It is necessary to delete standard SSM configuration in order to add group-list and it is necessary to delete all existing group-list configuration in order to configure standard SSM configuration.	

17.9.1.12 ip pim multipath next-hop

	ip pim [vrf <vrf-name>] multipath next-hop [<algorithm>] no ip pim [vrf <vrf-name>] multipath next-hop Configures PIM next-hop calculation algorithm. The no form of the command resets PIM next-hops configuration to default (highest neighbor).	
Syntax Description	vrf	VRF name
	algorithm	Selectable next-hop calculation algorithms: • g-hash - selects next-hop according to group address • mod - split groups between next hops on a module basis • s-g-hash - Selects next-hop according to group and source address
Default	Highest neighbor - next-hop with highest IP address is selected	
Configuration Mode	config	
History	3.6.8100	
	3.7.1100	Updated syntax
Example	`switch (config) # ip pim multipath next-hop g-hash`	
Related Commands		
Notes		

17.9.1.13 ip pim multipath rp

	ip pim multipath rp [<algorithm>] no ip pim multipath rp Configures PIM RP selection algorithm. The no form of the command resets PIM RP selection algorithm to default (g-hash algorithm which is described in RFC 4601, sec. 4.7.2).	
Syntax Description	algorithm	Selectable RP selection algorithms: • mod - split groups between RPs on a module basis
Default	G-hash—RPs are selected according to group address	
Configuration Mode	config	
History	3.7.1100	
Example	`switch (config) # ip pim multipath rp mod`	
Related Commands		
Note		

17.9.1.14 clear ip pim counters

	clear ip pim [vrf <vrf-name> | all] counters Clears PIM counter information.	
Syntax Description	vrf	VRF name or all VRFs
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6102	
Example	`switch (config) # clear ip pim counters`	
Related Commands		
Notes		

17.9.1.15 show ip pim protocol

	show ip pim [vrf {all | <vrf_name>}] protocol Displays PIM protocol information.	
Syntax Description	vrf	Displays output for a specific VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.6.6102	Updated example
	3.6.8008	Updated example and added “vrf” parameter
	3.7.1100	Updated description and example output
	3.9.1900	Updated example

Example
switch (config) # show ip pim vrf default protocol PIM Control Counters for VRF "default": PIM Mode: BIDIR Next-hop selection: highest neighbor RP selection: hash4601 (S,G) expiry timer: 210 seconds PIM Control Counters: --- Counters Received Sent Invalid --- Assert 0 0 0 Bootstrap Router 224 218 0 CRP Advertisement 0 0 0 Hello 443 551 0 J/P 0 0 0 Join 0 0 N/A Prune 0 0 N/A Register N/A N/A N/A Register Stop N/A N/A N/A State Refresh 0 0 0 DF Election 0 0 0

Related Commands	
Notes	

17.9.1.16 show ip pim bsr

	show ip pim [vrf {all | <vrf_name>}] bsr Displays PIM BSR information.	
Syntax Description	vrf	Displays output for a specific VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5006	
	3.6.6102	Updated example
	3.6.8008	Updated example and added “vrf” parameter
Example	switch (config) # show ip pim vrf all bsr PIMv2 Bootstrap information for VRF "default": No BSR is currently elected. This system is not a candidate-BSR PIMv2 Bootstrap information for VRF "vrf_1": BSR address : 17.17.17.10 Uptime : N/A BSR Priority : 64 Hash mask length : 30 Expires : 00:00:34 Candidate BSR : Yes Candidate BSR address: 17.17.17.10 priority : 64 hash mask length : 30 interval : N/A holdtime : N/A	
Related Commands		
Notes		

17.9.1.17 show ip pim interface

	show ip pim [vrf {all | <vrf_name>}] interface {[ethernet <port> | port-channel <id> | vlan <vlan id>]} Displays information about the enabled interfaces for PIM.	
Syntax Description	vrf	Displays output for a specific VRF
	ethernet <port>	Filters the output for specific Ethernet port
	port-channel <id>	Filters the output for specific LAG interface
	vlan <vlan-id>	Filters the output for specific VLAN interface
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5006	
	3.6.6102	Updated example
	3.6.8008	Updated example and added “vrf” parameter
	3.9.1900	Updated example

Example

```
switch (config)# show ip pim vrf default interface ethernet 1/17

VRF "default":
  Interface eth1/17 address 17.17.17.10:
    PIM                       : enabled
    PIM version               : 2
    PIM mode                  : bidir
    PIM DR                    : N/A
    PIM DR Priority           : N/A
    PIM configured DR priority: N/A
    PIM DF robustness         : 3
    PIM DF Offer interval     : 100 msec
    PIM DF Backoff interval   : 1000 msec
    PIM neighbor count        : 1
    PIM neighbor holdtime     : 105 secs
    PIM Hello Interval        : 30 seconds, next hello will be sent in: 00:00:00
    PIM Hello Generation ID   : d674dec2
    PIM Join-Prune Interval   : N/A
    PIM domain border         :

    PIM Interface Statistics:
      General (sent/received):
        Hellos     : 125 / 123
        JPs        : 7 / 164
        Asserts    : 0 / 0
        DF-Election: 1 / 2

      Errors:
        Checksum errors                        : N/A
        Invalid packet types/DF subtypes       : N/A / 0
        Authentication failed                  : N/A
        Packets from non-neighbors             : 0
        JPs received on RPF-interface          : N/A
        (*,G) Joins received with no/wrong RP  : N/A / N/A
        (*,G)/(S,G) JPs received for Bidir groups: 0
```

Related Commands	
Notes	

17.9.1.18 show ip pim interface brief

	show ip pim [vrf {all | <vrf_name>}] interface brief Displays PIM information summary for all interfaces.	
Syntax Description	vrf	Displays output for a specific VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5006	
	3.6.8008	Updated example and added “vrf” parameter
Example		

```
switch (config)# show ip pim vrf all interface brief

VRF "default":
  ------------------------------------------------------------------------
  Address         Interface     Ver/      Nbr    Query  DR     DR
                                Mode      Count  Intvl  Prior
  ------------------------------------------------------------------------
  20.20.20.10     eth1/1        v2/S      0      30     1      20.20.20.10
  30.30.30.10     eth1/2        v2/S      0      30     1      30.30.30.10
  17.17.17.10     eth1/17       v2/S      1      30     1      17.17.17.10
```

Related Commands	
Notes	

17.9.1.19 show ip pim neighbor

	show ip pim [vrf {all | <vrf_name>}] neighbor [vlan <vlan-id> | <other interfaces> | <ip-addr>] Displays information about IPv4 PIM neighbors.	
Syntax Description	vrf	Displays output for a specific VRF
	vlan <vlan-id>	Filters the output per specific VLAN ID
	neighbor-addr	Filters the output per specific neighbor IP address
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5006	
	3.6.8008	Updated example and added “vrf” parameter
	3.9.1900	Updated example: new option for mode Bidir or SM
Example		

```
switch (config) # show ip pim vrf default neighbor

VRF "default":
  ------------------------------------------------------------------------
  Neighbor      Interface     Uptime   Expires  Ver DR-Prio Mode      BFD
  ------------------------------------------------------------------------
  17.17.17.5    eth1/17       01:08:07 00:01:38 v2  N/A      Bidir    None
```

Related Commands	
Notes	

17.9.1.20 show ip pim rp

	show ip pim [vrf {all | <vrf_name>}] rp [<rp-address>] Displays information about the rendezvous points (RPs) for PIM.	
Syntax Description	vrf	Displays output for a specific VRF
	rp-address	Address of the rendezvous point
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5006	
	3.6.6102	Updated example
	3.6.8008	Updated example and added “vrf” parameter
	3.9.1900	Updated output: added PIM Bidir mode
Example	switch (config)# show ip pim vrf all rp PIM RP Status Information for VRF "default": PIM mode : BIDIR BSR : 100.100.100.100 expires : 53 priority : 64 hash-length: 30 RP 100.100.100.100: expires : 00:02:12 RP-source: 100.100.100.100 group ranges: 225.1.2.0/24, priority: 192 RP 100.100.100.100: expires : never RP-source: (local) group ranges: 224.0.0.0/4	
Related Commands		
Notes		

17.9.1.21 show ip pim rp-hash

	show ip pim [vrf <vrf-name> | all] rp-hash <group> Displays an RP that is selected for the given group.	
Syntax Description	vrf	VRF name of all VRFs
	group	A group address for RP calculation
Default	N/A	
Configuration Mode	Any command mode	

History	3.3.5006	
	3.7.1100	Updated example

Example	switch (config) # show ip pim rp-hash 224.1.1.0
	
	VRF "default":
	RP 192.167.7.1, v2:
	  RP-source:
	  priority : N/A
	  Uptime : N/A
	  Expires : N/A	

Related Commands		
Notes	RP is calculated according PIMv2 hash function as described in RFC 4601	

17.9.1.22 show ip pim rp-candidate

	show ip pim [vrf {all | <vrf_name>}] rp-candidate Displays information about RP candidate status.	
Syntax Description	vrf	Displays output for a specific VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5006	
	3.6.6000	Updated example
	3.6.6102	Updated example
	3.6.8008	Updated example and added “vrf” parameter
	3.9.1900	Updated example: added PIM mode (either BIDIR or SM)

Example	switch (config)# show ip pim vrf all rp-candidate
	
	VRF "default":
	  PIM mode: BIDIR
	
	VRF "vrf_1":
	  RP 17.17.17.10:
	  Interface : eth1/17
	  Interval : 60
	  Next advertisement in: 6
	  Holdtime : 150
	  Priority : 192
	
	  Group prefixes:
	  1: 225.0.0.0/24	

Related Commands		
Notes		

17.9.1.23 show ip pim ssm range

	show ip pim ssm [vrf {all | <vrf_name>}] range Displays information about configured PIM SSM ranges.	
Syntax Description	vrf	Displays information about configured PIM SSM ranges per specified VRF
Default	N/A	

Configuration Mode	Any command mode	
History	3.6.6000	
	3.6.6102	Updated example
	3.6.8008	Updated example and added “vrf” parameter

Example	switch (config)# show ip pim vrf all ssm range
	
	VRF "default":
	  PIM SSM is not configured
	
	VRF "vrf_1":
	  Range type : group-list
	  Total number of entries: 1
	
	  Group ranges:
	  1: 234.1.1.0/24
	  2: 234.1.2.0/24
	  3: 234.1.3.0/24
	  4: 234.1.4.0/24
	  5: 234.1.5.0/24	

Related Commands		
Notes		

17.9.1.24 show ip pim upstream joins

	show ip pim [vrf {all | <vrf_name>}] upstream joins Displays information about any PIM joins/prunes which are currently being sent to upstream PIM routers.	
Syntax Description	vrf	Displays output for a specific VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5006	
	3.6.6102	Updated example
	3.6.8008	Updated example and added “vrf” parameter

Example	switch (config) # show ip pim vrf all upstream joins
	
	VRF "default":
	  There are no upstream joins
	
	VRF "vrf_1":
	  Neighbor address 17.17.17.5:
	  via interface : 17.17.17.10
	  next message in: N/A seconds
	
	  Group 238.0.0.1:
	  Joins:
	  1: 10.10.10.5
	
	  Prunes:
	  No prunes included
	
	  Group 225.0.0.1:
	  Joins:
	  1: 10.10.10.5
	
	  Prunes:
	  No prunes included	

Related Commands		

Notes	Output contains the following information: neighbor address, interface address, group range, Joins, and Prunes.

17.9.2 PIM Bidir

17.9.2.1 ip pim bidir shutdown

	ip pim [vrf <vrf-name>] bidir shutdown [force] no ip pim [vrf <vrf-name>] bidir shutdown [force] **Disables** PIM bidirectional functionality, enabling PIM Sparse. The no form of the command **enables** PIM bidirectional functionality, disabling PIM Sparse.	
Syntax Description	vrf	VRF name.
	force	Keyword that is used in case a different mode already configured for PIM in the same VRF and some configuration is in place.
Default	Disabled for each VRF	
Configuration Mode	config	
History	3.9.1900	
Example	`switch (config) # ip pim [vrf <vrf-name>] bidir shutdown [force]` `switch (config) # no ip pim [vrf <vrf-name>] bidir shutdown [force]`	
Related Commands		

Notes	• If vrf <vrf-name> is not provided, the command will address vrf as “default”
	• When applying PIM mode BIDIR to vrf, the same mode will apply to ALL other VRFs with enabled PIM protocol
	• If a different mode already configured for PIM in the same VRF and "force" was not used, the following warning message will appear:
	“PIM SM configuration is present—please remove it to proceed or use “force” keyword to remove current configuration”
	• If another VRF PIM is enabled and is already configured with other PIM mode, force will not work and the warning message appear:
	“PIM SM configuration is present on other vrf—please remove it to proceed”
	• **Switch PIM mode to Sparse**
	The following commands are disabled in PIM Sparse mode:
	  • ip pim df-robustness
	  • ip pim df-offer-interval
	  • ip pim df-backoff-interval
	Remove configuration of the following (if **force** keyword provided):
	  • ip pim df-robustness
	  • ip pim df-offer-interval
	  • ip pim df-backoff-interval
	  • ip pim rp-candidate *
	  • ip pim rp-address *
	• **Switch PIM mode to Bidir**
	The following commands are disabled in PIM Bidir mode:
	  • ip pim register-source
	  • ip pim sg-expiry-timer
	  • ip pim ssm
	  • ip pim dr-priority
	  • ip pim join-prune-interval
	Remove configuration of the following (if **force** keyword provided):
	  • ip pim register-source
	  • ip pim sg-expiry-timer
	  • ip pim ssm
	  • ip pim dr-priority
	  • ip pim join-prune-interval
	  • ip pim rp-candidate *
	  • ip pim rp-address *

17.9.2.2 ip pim df-robustness

	ip pim df-robustness <number> no ip pim df-robustness Changes value of df-robustness. The no form of the command changes the value of df-robustness back to default	
Syntax Description	number	Value range: 1-255 Default: 3
Default	Disabled	
Configuration Mode	config bidir mode config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.9.1900	
Example	`switch (config interface vlan 10) # ip pim df-robustness <number>` `switch (config interface vlan 10) # no ip pim df-robustness`	
Related Commands		

Notes	The command “ip pim sparse-mode” must be run prior to using this command (available only in bidir mode) This command is part of the DF election mechanism: A router assumes the role of the DF after having advertised its metrics df-robustness times without receiving any offer from any other neighbor. At that point, it transmits a Winner message that declares to every other router on the link the identity of the winner and the metrics it is using. If a router hears a better offer than its own from a neighbor, it stops participating in the election for a period of (df-robustness * df-offer-interval), thus giving a chance to the neighbor with the better metric to be elected DF. If during this period no winner is elected, the router restarts the election from the beginning. If at any point during the initial election a router receives an out of order offer with worse metrics than its own, then it restarts the election from the beginning.

17.9.2.3 ip pim df-backoff-interval

	ip pim df-backoff-interval <milliseconds> no ip pim df-backoff-interval Changes the value of backoff interval. The no form of the command changes the value of backoff interval back to default.	
Syntax Description	milliseconds	Value range: 100-65,535 msec Default: 1000 msec
Default	Disabled	
Configuration Mode	config bidir mode config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.9.1900	
Example	`switch (config interface vlan 10) # ip pim df-backoff-interval <milliseconds>` `switch (config interface vlan 10) # no ip pim df-backoff-interval`	
Related Commands		
Notes	The command “ip pim sparse-mode” must be run prior to using this command (available only in bidir mode) This command is part of the DF election mechanism: Upon receipt of an offer that is better than its current metric, the DF records the identity and metrics of the offering router and responds with a Backoff message. This instructs the offering router to hold off for a short period of time while the unicast routing stabilizes, and other routers get a chance to put in their offers. The Backoff message includes the offering router's new metric and address. All routers on the link that have pending offers with metrics worse than those in the Backoff message (including the original offering router) will hold further offers for a period of time defined in the Backoff message. If a third router sends a better offer during the Backoff_Period, the Backoff message is repeated for the new offer and the Backoff_Period is restarted.	

17.9.2.4 ip pim df-offer-interval

	ip pim df-offer-interval <milliseconds> no ip pim df-offer-interval Changes value of Offer interval The no form of the command changes the value of Offer interval back to default

Syntax Description	milliseconds	Value range: 100-10,000 msec Default: 100 msec
Default	Disabled	
Configuration Mode	config bidir mode config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.9.1900	
Example	switch (config interface vlan 10) # ip pim df-offer-interval <milliseconds> switch (config interface vlan 10) # no ip pim df-offer-interval	
Related Commands		
Notes	The command “ip pim sparse-mode” must be run prior to using this command (available only in bidir mode) This command is part of the DF election mechanism: Initially, when no DF has been elected, routers finding out about a new RPA start participating in the election by sending Offer messages. Offer messages include the router's metric to reach the RPA(Rp-address). Offers are periodically retransmitted with a period of Offer_Interval.	

17.9.2.5 show ip pim interface df

	show ip pim [vrf {all | <vrf_name>}] interface {[ethernet <port> | port-channel <id> | vlan <vlan id>]} df Displays information about IPv4 PIM interface DF election per interface per RP.	
Syntax Description	Vrf <vrf-name>	If not provided will address vrf “default”
	Vrf all	Will show for all vrf
	ethernet <port>	Filters the output for specific Ethernet port
	port-channel <id>	Filters the output for specific LAG interface
	vlan <vlan-id>	Filters the output for specific VLAN interface
Default	Disabled for each VRF	
Configuration Mode	config bidir mode	
History	3.9.1900	

Example	show ip pim interface df
	
	VRF "default":
	
	--- 	
	-
	  Interface RP State DF Winner Metric
	Uptime
	
	--- 	
	-
	  eth1/12 100.100.100.100 Winner 2.1.1.2 0 0:29:48
	  eth1/12 100.100.100.100 Winner 2.1.1.2 0 0:29:48
	  eth1/13 100.100.100.100 Winner 1.1.1.2 0 0:29:43
	  eth1/13 100.100.100.100 Winner 1.1.1.2 0 0:29:43
	  eth1/3 100.100.100.100 Winner 192.168.2.254 0 5:10:42
	  eth1/3 100.100.100.100 Winner 192.168.2.254 0 5:10:42
	  eth1/5 100.100.100.100 Winner 192.168.3.254 0 5:10:42
	  eth1/5 100.100.100.100 Winner 192.168.3.254 0 5:10:43	

Related Commands		

Notes	The command “ip pim sparse-mode” must be run prior to using this command (available only in Bidir mode) This command is part of the DF election mechanism: Initially, when no DF has been elected, routers finding out about a new RPA start participating in the election by sending Offer messages. Offer messages include the router's metric to reach the RPA(RP-address). Offers are periodically retransmitted with a period of Offer_Interval. Table columns: Interface name RP: IP address of the RP for with this line is relevant State: state of DF-election: Winner\Loser\Disabled DF Winner: IP of the winner switch for this interface and RP Metric: metric towards the RP Uptime: uptime of the interface

17.9.3 Multicast

17.9.3.1 ip multicast-routing

	ip multicast-routing [vrf <vrf-name>] no ip multicast-routing [vrf <vrf-name>] Allows the switch to forward multicast packets. The no form of the command disables multicast routing.	
Syntax Description	vrf	VRF name
Default	Disabled	
Configuration Mode	config	
History	3.3.5006	
Example	`switch (config)# ip multicast-routing`	
Related Commands		
Notes		

17.9.3.2 ip mroute

	ip mroute [vrf <vrf-name>] {<ip-addr> <ip-mask> <next-hop>} [pref] no ip mroute [vrf <vrf-name>] {<ip-addr> <ip-mask> <next-hop>} Configure multicast reverse path forwarding (RPF) static routes. The no form of the command deletes the static multicast route.	
Syntax Description	ip-addr	Unicast IP address.
	ip-mask	Network mask in a dotted format (e.g. 255.255.255.0) or /24 format.
	next-hop	Next hop IP address.
	preference	Route preference. Range: 1-255.
Default	Preference is 1	
Configuration Mode	config	
History	3.3.5006	
	3.6.6000	Added “next-hop” parameter to “no” form

Example	`switch (config) # no ip mroute 2.1.1.0 /24 3.1.1.1`
Related Commands	
Notes	

17.9.3.3 ip multicast ttl-threshold

	ip multicast ttl-threshold <ttl-value> no ip multicast ttl-threshold Configures the time-to-live (TTL) threshold of packets being forwarded out of an interface. The no form of the command removes RPF static routes.	
Syntax Description	ttl-value	Range: 0-225
Default	0—all packets are forwarded	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.5006	
Example	`switch (config interface vlan 10)# ip multicast ttl-threshold 10`	
Related Commands		
Notes		

17.9.3.4 clear ip mroute

	clear ip mroute [vrf <vrf>] [<group-address> [<source-address>]] Clears multicast route information.	
Syntax Description	vrf	Clears multicast route information for specific VRF
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6102	
Example	`switch (config) # clear ip mroute 237.0.0.1 1.1.1.8`	
Related Commands		
Notes	This command does not support clearing specific (S,G) state if G belongs to an ASM group range. Here (S,G) refers to source and group parameters accordingly.	

17.9.3.5 show ip mroute

	show ip mroute [vrf {all | <vrf-name>}] [<group> [<prefix> [<source>]]] Displays information about IPv4 multicast routes.	
Syntax Description	source	Source IP address
	group	IP address of multicast group
	prefix	Network prefix of multicast group (in the format of /24, or 255.255.255.0 for example)

	summary	Displays a summary of the multicast routes
	vrf	Displays information pertinent to specified or all VRFs
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.1000	
	3.5.1000	Added new F flag and Updated example
	3.6.8008	Updated example and added “vrf” parameter
	3.8.1100	Added W/L line to Example output

Example

```
switch (config) # show ip mroute vrf vrf_1

IP Multicast Routing Table:
  Flags:
    B  : Bidir Group
    A  : ASM Group
    S  : SSM Group
    L  : Local
    P  : Pruned
    R  : RP-bit set
    T  : SPT-bit set
    J  : Join SPT
    F  : Failed to install in H/W
    W/L: Assert winner/loser

  Timers         : Uptime/Expires
  Interface state: Interface, State/Mode

VRF "vrf_1":
  (*, 225.0.0.1/32), 00D 00:04:40, RP 17.17.17.10, flags: AL:
    Incoming interface: eth1/17
    RPF Neighbor      : 0.0.0.0

    Outgoing interface list:
         eth1/1, N/A/ASM, 00D 00:04:40/00D 00:00:00

  (10.10.10.5, 225.0.0.1/32), 00D 00:04:37/00D 00:00:22, flags: AT:
    Incoming interface: eth1/17
    RPF Neighbor      : 17.17.17.5

    Outgoing interface list:
  (10.10.10.5, 225.0.0.2/32), 00D 00:04:31, flags: A:
    Incoming interface: eth1/17
    RPF Neighbor      : 17.17.17.5

    Outgoing interface list:
  (10.10.10.5, 225.0.0.3/32), 00D 00:04:16, flags: A:
    Incoming interface: eth1/17
    RPF Neighbor      : 17.17.17.5

    Outgoing interface list:
  (10.10.10.5, 238.0.0.1/32), 00D 00:04:40/00D 00:00:19, flags: ST:
    Incoming interface: eth1/17
    RPF Neighbor      : 17.17.17.5

    Outgoing interface list:
         eth1/2, N/A/SSM, 00D 00:04:40/00D 00:00:00


switch (config) # show ip mroute vrf vrf_1 225.0.0.1

IP Multicast Routing Table:
  Flags:
    B  : Bidir Group
    A  : ASM Group
    S  : SSM Group
    L  : Local
    P  : Pruned
    R  : RP-bit set
    T  : SPT-bit set
```

```
    J   : Join SPT
    F   : Failed to install in H/W
    W/L:

  Timers         : Uptime/Expires
  Interface state: Interface, State/Mode

VRF "vrf_1":
  (*, 225.0.0.1/32), 00D 00:13:27, RP 17.17.17.10, flags: AL:
    Incoming interface: eth1/17
    RPF Neighbor      : 0.0.0.0

    Outgoing interface list:
         eth1/1, N/A/ASM, 00D 00:13:27/00D 00:00:00

  (10.10.10.5, 225.0.0.1/32), 00D 00:13:24/00D 00:00:35, flags: AT:
    Incoming interface: eth1/17
    RPF Neighbor      : 17.17.17.5

    Outgoing interface list:


switch (config) # show ip mroute vrf all 225.0.0.1 /32

IP Multicast Routing Table:
  Flags:
    B   : Bidir Group
    A   : ASM Group
    S   : SSM Group
    L   : Local
    P   : Pruned
    R   : RP-bit set
    T   : SPT-bit set
    J   : Join SPT
    F   : Failed to install in H/W

    W/L:

  Timers         : Uptime/Expires
  Interface state: Interface, State/Mode

VRF "vrf_1":
  (*, 225.0.0.1/32), 00D 00:14:54, RP 17.17.17.10, flags: AL:
    Incoming interface: eth1/17
    RPF Neighbor      : 0.0.0.0

    Outgoing interface list:
         eth1/1, N/A/ASM, 00D 00:14:54/00D 00:00:00

  (10.10.10.5, 225.0.0.1/32), 00D 00:14:51/00D 00:00:08, flags: AT:
    Incoming interface: eth1/17
    RPF Neighbor      : 17.17.17.5

    Outgoing interface list:
```

Related Commands	
Notes	

17.9.3.6 show ip mroute summary

	show ip mroute [vrf {all | <vrf-name>}] summary Displays a summary of the IPv4 multicast routes.	
Syntax Description	vrf	Displays information pertinent to specified or all VRFs
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.1000	
	3.6.8008	Updated example

	3.8.1100	Added W/L line to Example output

Example

```
switch (config) # show ip mroute vrf vrf_1 summary

IP Multicast Routing Table:
  Flags:
    B  : Bidir Group
    A  : ASM Group
    S  : SSM Group
    L  : Local
    P  : Pruned
    R  : RP-bit set
    T  : SPT-bit set
    J  : Join SPT
    F  : Failed to install in H/W
    W/L:
  Timers         : Uptime/Expires
  Interface state: Interface, Next-Hop or VCD, State/Mode

VRF "vrf_1":
  (*, 225.0.0.1/32):
    Uptime   : 00D 00:11:18
    RP       : 17.17.17.10
    OIF count: 1
    flags    : AL

  (10.10.10.5, 225.0.0.1/32):
    Uptime   : 00D 00:11:15
    Exptime  : 00D 00:00:44
    OIF count: 0
    flags    : AT

  (10.10.10.5, 238.0.0.1/32):
    Uptime   : 00D 00:11:18
    Exptime  : 00D 00:00:41
    OIF count: 1
    flags    : ST

  Total: 3 routes
```

Related Commands	
Notes	

17.9.4 IGMP

17.9.4.1 ip igmp immediate-leave

	ip igmp immediate-leave no ip igmp immediate-leave Enables the device to remove the group entry from the multicast routing table immediately upon receiving a leave message for the group. The no form of the command disables immediate-leave.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface vlan config interface ethernet configured as a router port interface config interface port-channel configured as a router port interface
History	3.6.8100
Example	`switch (config interface vlan 10)# ip igmp immediate-leave`

Related Commands	
Notes	

17.9.4.2 ip igmp last-member-query-response-time

	ip igmp last-member-query-response-time <interval> no ip igmp last-member-query-response-time Configures the IGMP last member query response time in seconds. The no form of the command resets this parameter to its default.	
Syntax Description	interval	IGMP last member query response time. Range:1-25 seconds.
Default	1	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.5006	
Example	`switch (config interface vlan 10)# ip igmp last-member-query-response-time 10`	
Related Commands		
Notes	When both “IGMP” and “IGMP Snooping” handle a Leave message and have different values for “Last Member Query Time” timer configured, then traffic loss may occur for a short period of time.	

17.9.4.3 ip igmp startup-query-count

	ip igmp startup-query-count <count> no ip startup-query-count Configures the number of query messages an interface sends during startup. The no form of the command resets this parameter to its default.	
Syntax Description	count	Range: 1-255
Default	2	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.5006	
Example	`switch (config interface vlan 10)# ip igmp startup-query-count 10`	
Related Commands		
Notes		

17.9.4.4 ip igmp startup-query-interval

	ip igmp startup-query-interval <interval> no ip startup-query-interval Configures the IGMP startup query interval in seconds. The no form of the command resets this parameter to its default.

Syntax Description	interval	Range: 1-1800 seconds
Default	31	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.5006	
Example	`switch (config interface vlan 10)# ip igmp startup-query-interval 10`	
Related Commands		
Notes		

17.9.4.5 ip igmp query-interval

	ip igmp query-interval <interval> no ip igmp query-interval Configures the IGMP query interval in seconds. The no form of the command resets this parameter to its default.	
Syntax Description	interval	The IGMP query interval Range: 1-1800 seconds
Default	125 seconds	
Configuration Mode	config interface vlan	
History	3.3.5006	
Example	`switch (config interface vlan 10)# ip igmp query-interval 60`	
Related Commands		
Notes		

17.9.4.6 ip igmp query-max-response-time

	ip igmp query-max-response-time <time> no ip igmp query-max-response-time Configures the IGMP max response time in seconds. The no form of the command resets this parameter to its default.	
Syntax Description	time	The IGMP max response time Range: 1-25 seconds
Default	10	
Configuration Mode	config interface vlan	
History	3.3.5006	
Example	`switch (config interface vlan 10)# ip igmp query-max-response-time 20`	
Related Commands		
Notes		

17.9.4.7 ip igmp robustness-variable

	ip igmp robustness-variable <count> no ip igmp robustness-variable Configures the IGMP robustness variable. The no form of the command resets this parameter to its default.	
Syntax Description	count	IGMP robustness variable Range: 1-7
Default	2	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.5006	
Example	`switch (config interface vlan 10)# ip igmp robustness-variable 4`	
Related Commands		
Notes	The robustness variable can be increased to increase the number of times that packets are resent This parameter reflects expected packet loss on a congested network	

17.9.4.8 ip igmp static-oif

	ip igmp static-oif <group> [source-ip <address>] no ip igmp static-oif <group> [source-ip <address>] Statically binds an IP interface to a multicast group. The no form of the command deletes the static multicast address from the interface.	
Syntax Description	group	Multicast IP address
	source-ip	IP address from which to receive group traffic
Default	N/A	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.5006	
Example	`switch (config interface vlan 10)# ip igmp static-oif 10.10.10.5`	
Related Commands		
Notes	PIM must be enabled in order to configure the route in the hardware.	

17.9.4.9 clear ip igmp groups

	clear ip igmp groups {all | interface <if> | vrf <number> | <group-address> <mask>} Clears IGMP group information.	
Syntax Description	all	Clears all IGMP groups
	interface	Clears IGMP groups on specific interface
	vrf	Clears IGMP groups in specific VRF

	group-address	Clears a specific group range
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
Example	`switch (config)# clear ip igmp groups all`	
Related Commands		
Notes		

17.9.4.10 show ip igmp groups

	show ip igmp [vrf {all |<vrf_name>}] groups [<group> | <iface>] Displays information about IGMP-attached group membership.	
Syntax Description	vrf	Displays output for a specific VRF
	group	Filters the output to a specific IP multicast group address
	iface	Filters the output to a specific IP interface (i.e. ethernet, port-channel, vlan interface)
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.6.6102	Updated example
	3.6.8008	Updated example and added “vrf” parameter
Example		

```
switch (config)# show ip igmp vrf all groups

IGMP Connected Group Membership
Type: S - Static, D - Dynamic

VRF "default":
  No IGMP group memberships learned or configured

VRF "vrf_1":
  ------------------------------------------------------------------------------------------
  Group Address         Type         Interface         Uptime         Expires         Last Reporter
  ------------------------------------------------------------------------------------------
  225.0.0.1             D            eth1/1            01:03:03       00:03:51        20.20.20.5
  238.0.0.1             D            eth1/2            01:03:03       N/A             30.30.30.5
```

Related Commands	
Notes	

17.9.4.11 show ip igmp interface

	show ip igmp [vrf <vrf-name> | all] interface [ethernet <if> | port-channel <if> | vlan <vlanid>] brief Displays IGMP brief configuration and status.	
Syntax Description	vrf	Displays output for a specific VRF
	brief	Displays brief output information

	ethernet	Displays output for a specific Ethernet port
	port-channel	Displays output for a specific LAG
	vlan <vlan-id>	Displays output for a specific VLAN ID
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.6.6102	Updated example
	3.6.8008	Updated example and added “vrf” parameter
	3.6.8100	Added “IGMP interface immediate leave” line to output

Example

```
switch (config)# show ip igmp interface vlan 10

Interface vlan10
        Status: protocol-down/link-down/admin-up
        VRF: "vrf-default"
        IP address: 10.10.10.1/24
        Active querier: 10.10.10.1
        Version: 2
        Next query will be sent in: 00:01:45
        Membership count: 0
        IGMP version: 2
        IGMP query interval: 125 secs
        IGMP max response time: 10 secs
        IGMP startup query interval: 31 secs
        IGMP startup query count: 2
        IGMP last member query interval: 1 secs
        IGMP last member query count: 2
        IGMP group timeout: 260 secs
        IGMP querier timeout: 0 secs
        IGMP unsolicited report interval: 10 secs
        IGMP robustness variable: 2
        IGMP interface immediate leave: Disabled
        Multicast routing status on interface: Enabled
        Multicast TTL threshold: 0

        IGMP interface statistics:
          General (sent/received):
            v2-queries: 2/0
            v2-reports: 0/0
            v2-leaves : 0/0
            v3-queries: 0/0
            v3-reports: 0/0

          Errors:
            Checksum errors                         : 0
            Packet length errors                    : 0
            Packets with Local IP as source         : 0
            Source subnet check failures            : 0
            Query from non-querier                  : 0
            Report version mismatch                 : 0
            Query version mismatch                  : 0
            Unknown IGMP message type               : 0
            Invalid v2 reports                      : 0
            Invalid v3 reports                      : 0
            Invalid leaves                          : 0
            Packets dropped due to router-alert check: 0
```

Related Commands	
Notes	

17.9.4.12 show ip igmp interface brief

	show ip igmp interface [ethernet <if> | port-channel <if> | vlan <vlan-id>] brief Displays brief IGMP configuration and status information.	
Syntax Description	vrf	Displays output for a specific VRF
	ethernet	Displays output for a specific Ethernet port
	port-channel	Displays output for a specific LAG
	vlan <vlan-id>	Displays output for a specific VLAN ID
Default	N/A	
Configuration Mode	Any command mode	
History	3.3.5200	
	3.6.6102	Updated example
	3.6.8008	Updated example and added “vrf” parameter
Example		

```
switch (config)# show ip igmp vrf all interface brief

VRF "default":
  ------------------------------------------------------------------------------------------
  Interface           IP Address           IGMP Querier         Membership Count     Version
  ------------------------------------------------------------------------------------------
  eth1/10             12.14.192.5          0.0.0.0              0                    v3

VRF "vrf_1":
  ------------------------------------------------------------------------------------------
  Interface           IP Address           IGMP Querier         Membership Count     Version
  ------------------------------------------------------------------------------------------
  eth1/1              20.20.20.10          20.20.20.10          1                    v2
  eth1/2              30.30.30.10          30.30.30.10          1                    v3
  eth1/17             17.17.17.10          17.17.17.5           0                    v3
```

Related Commands	
Notes	

17.10 IGMP Snooping

The Internet Group Multicast Protocol (IGMP) is a communications protocol used by hosts and adjacent routers on IP networks to establish multicast group memberships. The host joins a multicast-group by sending a join request message towards the network router, and responds to queries sent from the network router by dispatching a join report.

A given port can be either manually configured to be a MRouter port or it can be dynamically manifested when having received a query, hence, the network router is connected to this port. All IGMP Snooping control packets received from hosts (joins/leaves) are forwarded to the MRouter port, and the MRouter port updates its multicast-group database accordingly. Each dynamically learned multicast group will be added to all of the MRouter ports on the switch.

As many as 5K multicast groups can be created on the switch.

17.10.1 Configuring IGMP Snooping

IGMP snooping can be configured to establish multicast group memberships.

1. Enable IGMP snooping globally. Run:

```
switch (config) # ip igmp snooping
```

2. Enable IGMP snooping on a VLAN. Run:

```
switch (config) # vlan 2
switch (config vlan 2) # ip igmp snooping
```

17.10.2 Defining a Multicast Router Port on a VLAN

A Multicast Router (MRouter) port on can be defined on a VLAN in one of the methods described below:

- To change the interface switchport to trunk:
  a. Enable IGMP snooping globally. Run:

```
switch (config) # ip igmp snooping
```

  b. Change the interface switchport mode of the port (the interface is member of VLAN 1 by default). Run:

```
switch (config) # interface ethernet 1/1
switch (config interface ethernet 1/1) # switchport mode trunk
```

  c. Change back to config mode. Run:

```
switch (config interface ethernet 1/1) # exit
switch (config) #
```

  d. Define the MRouter port on the VLAN. Run:

```
switch (config) # vlan 2
switch (config vlan 2) # ip igmp snooping mrouter interface ethernet 1/1
```

- To change the interface switchport to hybrid:
  a. Enable IGMP snooping globally. Run:

```
switch (config) # ip igmp snooping
```

  b. Create a VLAN. Run:

```
switch (config) # vlan 200
switch (config vlan 200) #
```

  c. Change back to config mode. Run:

```
switch (config vlan 200) # exit
switch (config) #
```

d. Change the interface switchport mode of the port (the interface is member of VLAN 1 by default). Run:

```
switch (config) # interface ethernet 1/22
switch (config interface ethernet 1/22) # switchport mode hybrid
```

e. Attach the VLAN to the port's interface. Run:

```
switch (config interface ethernet 1/22) # switchport mode hybrid allowed-vlan 200
switch (config interface ethernet 1/22) #
```

f. Change to config mode again. Run:

```
switch (config interface ethernet 1/22) # exit
switch (config) #
```

g. Define the MRouter port on the VLAN. Run:

```
switch (config) # vlan 200
switch (config vlan 200) # ip igmp mrouter interface ethernet 1/22
```

- To change the interface switchport to access:
  a. Enable IGMP snooping globally. Run:

```
switch (config) # ip igmp snooping
```

  b. Create a VLAN. Run:

```
switch (config) # vlan 200
switch (config vlan 200) #
```

  c. Change back to config mode. Run:

```
switch (config vlan 200) # exit
switch (config) #
```

  d. Change the interface switchport mode of the port (the interface is member of VLAN 1 by default). Run:

```
switch (config) # interface ethernet 1/22
switch (config interface ethernet 1/22) # switchport mode access
```

  e. Attach the VLAN to the port's interface. Run:

```
switch (config interface ethernet 1/22) # switchport access vlan 200
```

  f. Change to config mode again. Run:

```
switch (config interface ethernet 1/22) # exit
```

  g. Define the MRouter port on the VLAN. Run:

```
switch (config) # vlan 200
switch (config vlan 200) # ip igmp mrouter interface ethernet 1/22
```

17.10.3 IGMP Snooping Querier

IGMP Snooping Querier complements the IGMP snooping functionality. IGMP Snooping Querier is used to support IGMP snooping in a VLAN where PIM and IGMP are not configured because the multicast traffic does not need to be routed. When IGMP Snooping Querier is enabled, IGMP queries are sent out periodically by the switch through all ports in the VLAN and to which hosts wishing to receive IP multicast traffic respond with IGMP report messages. IGMP Snooping Querier must be used in conjunction with IGMP snooping as IGMP snooping listens to these IGMP reports to establish appropriate forwarding.

To configure IGMP Snooping Querier:

1. Enable the IGMP snooping on the switch. Run:

```
switch (config) # ip igmp snooping
```

2. Create a VLAN and enable IGMP Snooping on VLAN. Run:

```
switch (config) # vlan 10
switch (config vlan 10)# ip igmp snooping
```

3. Enable the IGMP snooping querier on a specific VLAN. Run:

```
switch (config vlan 10)# ip igmp snooping querier
```

4. Set the query interval time. Run:

```
switch (config vlan 10)# ip igmp snooping querier query-interval 100
```

5. (Optional) Verify the IGMP snooping querier configuration. Run:

```
switch (config vlan 10)# show ip igmp snooping querier
Snooping querier information for VLAN 10

IGMP Querier Present
Querier IP address: 1.1.1.2
Query interval: 125
Response interval: 100
Group membership interval: 1
Robustness: 2
Version: 2
```

17.10.4 IGMP Snooping Querier Guard

In some environments, devices attached to a switch (such as hosts or other switches) cannot be managed by the switch administrator. This can lead to IGMP resources misconfiguration or abuse and is an operational behavior and security concern.
This is common in shared network infrastructures, where a 3rd party is connected to the switch to access resources that are made available via that network device.

IGMP Snooping Querier Guard enables the switch administrator to define a filter to discard IGMP Membership Query messages, allowing it to be selected as the IGMP querier by ignoring the received messages. Connecting a device to an interface where this filter is defined stops the IGMP Querier election process that allows a 3rd party device to trigger the local interface to be demoted from being the IGMP querier.

IGMP Snooping Querier Guard can be configured on specific interfaces such as a port, MLAG port channel, or port channel. It only works when "igmp snooping" is enabled.

To configure IGMP Snooping Querier Guard on a specific interface, do the following:

1. Enable the IGMP snooping on the switch. Run:

```
switch (config) # ip igmp snooping
```

2. Enable IGMP snooping querier-guard on a specific interface. Run:

```
switch (config interface ethernet 1/1) # ip igmp snooping querier-guard
```

17.10.5 IGMP Snooping Commands

- 17.10.5.1 ip igmp snooping (admin)
- 17.10.5.2 ip igmp snooping (config)
- 17.10.5.3 ip igmp snooping fast-leave
- 17.10.5.4 ip igmp snooping mrouter
- 17.10.5.5 ip igmp snooping static-group
- 17.10.5.6 ip igmp snooping querier
- 17.10.5.7 ip igmp snooping querier-guard
- 17.10.5.8 ip igmp snooping querier address
- 17.10.5.9 igmp snooping querier query-interval
- 17.10.5.10 ip igmp snooping profile
- 17.10.5.11 ip igmp snooping filter profile
- 17.10.5.12 ip igmp snooping max-groups
- 17.10.5.13 ip igmp version
- 17.10.5.14 clear ip igmp snooping counters
- 17.10.5.15 clear ip igmp snooping filter
- 17.10.5.16 show ip igmp snooping
- 17.10.5.17 show ip igmp snooping groups
- 17.10.5.18 show ip igmp snooping interfaces
- 17.10.5.19 show ip igmp snooping membership
- 17.10.5.20 show ip igmp snooping mrouter
- 17.10.5.21 show ip igmp snooping querier
- 17.10.5.22 show ip igmp snooping querier-guard
- 17.10.5.23 show ip igmp snooping querier counters
- 17.10.5.24 show ip igmp snooping statistics
- 17.10.5.25 show ip igmp snooping vlan
- 17.10.5.26 show ip igmp snooping profile
- 17.10.5.27 show ip igmp snooping filter

17.10.5.1 ip igmp snooping (admin)

	ip igmp snooping no ip igmp snooping Enables IGMP snooping globally or per VLAN. The no form of the command disables IGMP snooping globally or per VLAN.
Syntax Description	N/A
Default	IGMP snooping is disabled globally and per VLAN
Configuration Mode	config config vlan
History	3.1.1400

Example	`switch (config) # ip igmp snooping`
	
	`switch (config vlan 10) # ip igmp snooping`

Related Commands	show ip igmp snooping
Notes	IGMP snooping has global admin state, and per VLAN admin state. Both states need to be enabled in order to enable the IGMP snooping on a specific VLAN.

17.10.5.2 ip igmp snooping (config)

	ip igmp snooping {last-member-query-interval <1-25> | proxy reporting mrouter-timeout <60-600> | port-purge-timeout <130-1225> | report-suppression-interval <1-25> | unregistered multicast {flood | forward-to-mrouter-ports} | version {2 | 3}} no ip igmp snooping {last-member-query-interval | proxy reporting | mrouter-timeout | report-suppression-interval | unregistered multicast | version} Configures global IGMP parameters. The no form of the command resets the global IGMP parameters to default.	
Syntax Description	last-member-query-interval <1-25>	Sets the time period (in seconds) with which the general queries are sent by the IGMP querier. After timeout expiration, the port is removed from the multicast group.
	proxy reporting	Enables proxy reporting
	mrouter-timeout <60-600>	Sets the IGMP snooping MRouter port purge time-out after which the port gets deleted if no IGMP router control packets are received
	port-purge-timeout <130-1225>	Sets the IGMP snooping port purge time interval after which the port gets deleted if no IGMP reports are received
	report-suppression-interval <1-25>	Sets the IGMP snooping report-suppression time interval for which the IGMPv2 report messages for the same group will not get forwarded onto the MRouter ports
	unregistered multicast	Sets the behavior of the snooping switch for unregistered multicast traffic • flood - flood unregistered multicast traffic on all port in specific VLAN • forward-to-mrouter-ports - forward unregistered multicast traffic only to mrouter ports in specific VLAN

	version	Sets the default operating version to use for newly created IGMP snooping instances • 2 - enables IGMPv2 • 3 - enables IGMPv3 Also available in “config vlan” configuration mode
Default	last-member-query-interval - 1 second proxy reporting - disabled mrouter-timout - 125 port-purge-timeout - 260 seconds report-suppression-interval - 5 seconds unregistered multicast - flood version - 3	
Configuration Mode	config	
History	3.1.1400	
	3.2.0500	Added “unregistered multicast” parameter
	3.6.1002	Added “version parameter”
	3.6.2100	Changed default value for “version” parameter
	3.7.1100	Updated note
Example	`switch (config) # ip igmp snooping report-suppression-interval 3`	
Related Commands	ip igmp snooping (admin) show ip igmp snooping	
Notes	When both IGMP and IGMP snooping protocols handle a Leave message and have different values for “Last Member Query Time” timer configured, then there is traffic loss for a short period of time.	

17.10.5.3 ip igmp snooping fast-leave

	ip igmp snooping fast-leave no ip igmp snooping fast-leave Enables fast leave processing on a specific interface. The no form of the command disables fast leave processing on a specific interface.	
Syntax Description	N/A	
Default	Enabled	
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel	
History	3.1.1400	
	3.3.4500	Added MPO configuration mode
Example	`switch (config interface ethernet 1/1) # ip igmp snooping fast-leave`	
Related Commands	show ip igmp snooping interfaces	
Notes		

17.10.5.4 ip igmp snooping mrouter

	ip igmp snooping mrouter interface <type> <number> no ip igmp snooping mrouter interface <type> <number> Creates a static multicast router port on a specific VLAN, on a specific interface. The no form of the command removes the static multicast router port from a specific VLAN.	
Syntax Description	interface <type> <number>	Attaches the group to a specific interface type - ethernet or port-channel
Default	No static mrouters are configured	
Configuration Mode	config vlan	
History	3.1.1400	
Example	`switch (config vlan 1) # ip igmp snooping mrouter 1/1`	
Related Commands	show ip igmp snooping mrouter	
Notes	The multicast router port can be created only if IGMP snooping is enabled both globally and on the VLAN.	

17.10.5.5 ip igmp snooping static-group

	ip igmp snooping static-group <IP address> interface <type> <number> [source <source-ip>] no ip igmp snooping static-group <IP address> interface <type> <number> [source <source-ip>] Creates a specified static multicast group for specified ports and from a specified source IP address. The no form of the command deletes the interface from the multicast group.	
Syntax Description	IP address	Multicast IP address <224.x.x.x - 239.255.255.255>
	interface	Attach the group to a specific interface
	type	Ethernet or port-channel
	source	Source IP address. If omitted, a multicast group is created for all sources.
Default	No static groups are configured	
Configuration Mode	config vlan	
History	3.1.1400	
	3.6.2100	Added “source” parameter
Example	`switch (config vlan 1) # ip igmp snooping static-group 230.0.0.1 1/1`	
Related Commands	show ip igmp snooping groups	
Notes	If the deleted interface is the last port, it deletes the entire multicast group.	

17.10.5.6 ip igmp snooping querier

	ip igmp snooping querier no ip igmp snooping querier Enables the IGMP Snooping Querier on a VLAN. The no form of the command disables the IGMP Snooping Querier on a VLAN.
Syntax Description	N/A
Default	Disable
Configuration Mode	config vlan
History	3.3.4200
Example	`switch (config vlan 1)# ip igmp snooping querier`
Related Commands	igmp snooping querier query-interval show ip igmp snooping querier
Notes	

17.10.5.7 ip igmp snooping querier-guard

	ip igmp snooping querier-guard no ip igmp snooping querier-guard Enables IGMP querier guard functionality on per L2 interface basis. The no form of the command disables IGMP querier guard functionality on the current interface.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config interface ethernet config interface port-channel config interface mlag-port-channel
History	3.8.2000
Example	`switch (config interface ethernet 1/1) # ip igmp snooping querier-guard`
Related Commands	show ip igmp snooping querier-guard show ip igmp snooping interfaces
Notes	Doesn't affect layer 3 multicast router.

17.10.5.8 ip igmp snooping querier address

	ip igmp snooping querier address <ip_address> ip igmp snooping querier Configures the IGMP Snooping querier source IP address. The no form of the command deletes the querier IP address.	
Syntax Description	ip_address	The querier IP address
Default	Disabled	
Configuration Mode	config vlan	
History	3.4.2000	

Example	`switch (config vlan 1) # ip igmp snooping querier address 1.1.1.2`
Related Commands	ip igmp snooping querier ip igmp snooping querier query-interval show ip igmp snooping querier
Notes	Need to configure the querier IP address, otherwise the "0.0.0.0." address will be used.

17.10.5.9 igmp snooping querier query-interval

	igmp snooping querier query-interval <time> no igmp snooping querier query-interval Configures the query interval. The no form of the command rests the parameter to its default.	
Syntax Description	time	Time interval between queries (in seconds).
Default	125 seconds	
Configuration Mode	config vlan	
History	3.3.4200	
	3.7.1000	Updated example
Example	`switch (config vlan 1)# igmp snooping querier query-interval 100`	
Related Commands	igmp snooping querier query-interval show ip igmp snooping querier	
Notes		

17.10.5.10 ip igmp snooping profile

	ip igmp snooping profile <profile_name> [seq <num>]{permit|deny} {group_address[/prefix_length]} [source_address[/prefix_length]] no ip igmp snooping profile <profile_name> [seq <num>] Defines an IGMP Snooping Filter Profile and rules of the IGMP Snooping Filter Profile. The no form of the command deletes the profile and the rules.	
Syntax Description	profile_name	User specified profile name.
	seq <number>	Sequence number: 1-65534.
	permit	Permits access for a matching condition.
	deny	Denies access for a matching condition.
	group_address[/prefix_length]	Group IP address or prefix.
	source_address[/prefix_length]	Source IP address or prefix.
Default	Sequence value: 10	
Configuration Mode	config	
History	3.9.2100	

Example	switch (config)# ip igmp snooping profile proflie_1 switch (config ip igmp snooping profile proflie_1)# permit 224.1.1.0/24 192.168.1.1 switch (config ip igmp snooping profile proflie_1)# deny 224.1.1.0/24 192.168.1.1 switch (config ip igmp snooping profile proflie_1)# seq 53 permit 224.2.0.1 192.168.53.0/24 switch (config ip igmp snooping profile proflie_1)# seq 54 permit 224.3.0.0/16 192.168.54.1
Related Commands	show ip igmp snooping profile
Notes	• By default, rules sequence numbers are incremented decimally (i.e., 10, 20, 30, 40). • Up to 32 user defined rules per profile are permitted. • There is always a silent “deny any” rule with seq number 65535 at the end of each profile rule list. • Group prefix, source prefix defined in rules are applied to filter only those group address and source address list inside the incoming IGMP snooping reports, not considering other attributes (e.g., record type EXCLUDE, INCLUDE, and so forth).

17.10.5.11 ip igmp snooping filter profile

	interface {ethernet <port>[-<port>] | port-channel <lag-id>[-<lag-id>] | mlag-port-channel <mlag-id>[-<mlag-id>]}ip igmp snooping filter profile <profile_name> [vlan <num>[-<range>]] no interface {ethernet <port>[-<port>] | port-channel <lag-id>[-<lag-id>] | mlag-port-channel <mlag-id>[-<mlag-id>]}ip igmp snooping filter profile <profile_name> [vlan <num>[-<range>]] Applies a defined IGMP snooping filter profile to an interface and corresponding VLANs.	
Syntax Description	port	Ethernet port
	lag-id	LAG ID
	mlag-id	MLAG ID
	profile_name	Specified name
	vlan <num>[-<range>]	Specified VLAN or specified VLAN range
Default	N/A	
Configuration Mode	config	
History	3.9.2100	
Example	switch (config)# interface ethernet 1/21 ip igmp snooping filter profile WEB-Profile vlan 1	
Related Commands	show ip igmp snooping profile	
Notes	• If “vlan” parameter is not defined, the command will apply to all VLANs • IGMP filtering takes effect on the local switch only. After filtering, the original packet is not directly discarded and continues to be forwarded to uplink switches. That is, the forwarding packet is not the filtered/changed packet (or remaining part of the packet), but the original/unchanged packet itself.	

17.10.5.12 ip igmp snooping max-groups

	interface {ethernet <port>[-<port>] | port-channel <lag-id>[-<lag-id>] | mlag-port-channel <mlag-id>[-<mlag-id>]} ip igmp snooping max-groups <value> no interface {ethernet <port>[-<port>] | port-channel <lag-id>[-<lag-id>] | mlag-port-channel <mlag-id>[-<mlag-id>]} ip igmp snooping max-groups <value> vlan <num>[-<range>] ip igmp snooping max-groups <value> no vlan <num>[-<range>] ip igmp snooping max-groups Applies maximum number of IGMP groups that can be joined on a specific interface or in a specific VLAN. The no form of the command cancels the maximum number of IGMP groups that can be joined on a specific interface or in a specific VLAN.	
Syntax Description	port	Ethernet port
	lag-id	LAG ID
	mlag-id	MLAG ID
	vlan <num>[-<range>]	Specified VLAN or specified VLAN range
	max-groups <value>	Maximum number of IGMP groups Range: 1-32767
Default	N/A	
Configuration Mode	config	
History	3.9.2100	

Example	`switch (config) # interface ethernet 1/2 ip igmp snooping max-groups 100`
	`switch (config) # vlan 10 ip igmp snooping max-groups 100`
	`switch (config) # no  interface ethernet 1/21 ip igmp snooping max-groups`	

Related Commands	show ip igmp snooping profile	
Notes	• For existing groups registered before enabling max-group filtering, a report packet is accepted in order to refresh the existing groups • IGMP filtering takes effect on the local switch only. After filtering, the original packet is not directly discarded and continues to be forwarded to uplink switches. That is, the forwarding packet is not the filtered/changed packet (or remaining part of the packet), but the original/unchanged packet itself.	

17.10.5.13 ip igmp version

	ip igmp version <2, 3> no ip igmp version Sets IGMP version on interface. The no form of the command resets the IGMP version on the interface to default value.	
Syntax Description	version	Protocol IGMP version. Range: 2-3
Default	IGMP version 2	
Configuration Mode	config interface vlan config interface ethernet (configured as a router port interface) config interface port-channel (configured as a router port interface)	
History	3.3.5006	
	3.8.1300	Added the command to the user manual

Example	`switch (config interface vlan 10) # ip igmp version 3`
Related Commands	
Notes	

17.10.5.14 clear ip igmp snooping counters

	clear ip igmp snooping counters [vlan <vlan-id>] Clears IGMP snooping counters.	
Syntax Description	vlan	Clears IGMP snooping counters per VLAN
Default	N/A	
Configuration Mode	config	
History	3.6.1002	
	3.6.6000	Updated command format
Example	`switch (config) # clear ip igmp snooping counters vlan 2`	
Related Commands		
Notes		

17.10.5.15 clear ip igmp snooping filter

	clear ip igmp snooping filter [interface {ethernet <port>[-<port>] | port-channel <lag-id>[-<lag-id>] | mlag-port-channel <mlag-id>[-<mlag-id>] } | vlan <num>] counters Clears the IGMP snooping filter counters for all interfaces or the specifically selected one(s).	
Syntax Description	port	Ethernet port
	lag-id	LAG ID
	mlag-id	MLAG ID
	vlan <num>	Specified VLAN
Default	N/A	
Configuration Mode	Any command mode	
History	3.9.2100	

Example	`switch (config) # clear ip igmp snooping filter counters`
	`switch (config) # clear ip igmp snooping filter interface ethernet 1/2 counters`
	`switch (config) # clear ip igmp snooping filter vlan 50 counters`	

Related Commands		
Notes		

17.10.5.16 show ip igmp snooping

	show ip igmp snooping Displays IGMP snooping information for all VLANs or a specific VLAN.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.1400	
	3.6.1002	Added default IGMP version to Example
	3.6.6102	Updated example

Example	`switch (config) # show ip igmp snooping`
	
	`IGMP snooping global configuration:`
	`IGMP snooping globally: enabled`
	`IGMP default version for new VLAN: V3`
	`IGMP snooping operationally: enabled`
	`Proxy-reporting globally: enabled`
	`Last member query interval: 1 seconds`
	`Mrouter timeout: 125 seconds`
	`Port purge timeout: 260 seconds`
	`Report suppression interval: 5 seconds`
	`IGMP snooping unregistered multicast: flood`	

Related Commands		
Notes		

17.10.5.17 show ip igmp snooping groups

	show ip igmp snooping groups [vlan <vid> [group <group-ip>]] Displays per VLAN the list of multicast groups attached (static or dynamic allocated) per port.	
Syntax Description	vid	VLAN ID
	group	Multicast group IP address
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.1400	
	3.6.1002	Updated example
	3.6.2100	Added “vlan” and “group” parameters and Updated example
	3.6.6102	Updated example output
Example		

```
switch (config) # show ip igmp snooping groups

------------------------------------------------
Vlan ID     Group           St/Dyn     Ports
------------------------------------------------
1           230.0.0.1       St         Eth1/1,Eth1/2
2           230.0.0.1       St         Eth1/4,Eth1/6
2           230.0.0.2       St         Eth1/5

Total Num of Dynamic Group Addresses: 1
Total Num of Static Group Addresses: 1

switch (config) # show ip igmp snooping groups vlan 1
--------------------------------------
Group           St/Dyn     Ports
--------------------------------------
230.0.0.1       St         Eth1/1,Eth1/2,Eth1/3

Total Num of Dynamic Group Addresses: 0
Total Num of Static Group Addresses: 1

switch (config) # show ip igmp snooping groups vlan 1 group 230.0.0.1
Snooping group information for VLAN 1 and group 230.0.0.1
        Filter Mode: EXCLUDE
        Exclude sources: None
        V1/V2 Receiver Ports: Eth1/1,Eth1/2,Eth1/3
        V3 Receiver Ports:   None
```

Related Commands	
Notes	

17.10.5.18 show ip igmp snooping interfaces

	show ip igmp snooping interfaces Displays IGMP snooping interface information.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.1400	
	3.8.2000	Updated example
	3.9.2000	Updated example
	3.9.2100	Updated example, adding support for IGMP snooping filtering

Example	`switch (config) # show ip igmp snooping interfaces`
	`interface  leave-mode  querier-guard  profile_filter  max-groups`
	`Eth1/1  Normal  Disabled  N/A  unlimited`
	`Eth1/2  Normal  Disabled  N/A  100`
	`Eth1/3  Normal  Enabled  profile_1  unlimited`
	`Eth1/4  Normal  Enabled  prof_2  200`
	`Eth1/5  Normal  Disabled  N/A  50`	

Related Commands	ip igmp snooping querier-guard ip igmp snooping fast-leave ip igmp snooping max-groups ip igmp snooping filter profile	
Notes	The "profile_filter" and "max-groups" columns are just placeholders for the IGMP Snooping filter feature that will be introduces 3.9.2100 or later.	

17.10.5.19 show ip igmp snooping membership

	show ip igmp snooping membership [vlan <vid> [group <group-ip>]] Displays information about host membership for multicast groups.	
Syntax Description	vlan	Displays IGMP snooping querier counters on specific VLAN
	group	Multicast group IP address
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.2100	
Example		

```
switch (config) # show ip igmp snooping membership vlan 1 group 224.5.5.5
Snooping membership information for VLAN 1 and group 224.5.5.5

Receiver Port: Eth1/1
  Attached Host: 10.10.10.1
    Version: 3
    Mode: Include
    Sources: 10.10.10.100
    Timeout since the host has been joined: 0:00:02
    Expiry timeout: 0:04:18
```

Related Commands	
Notes	

17.10.5.20 show ip igmp snooping mrouter

	show ip igmp snooping mrouter Displays IGMP snooping multicast router information.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.1400

Example	`switch (config) # show ip igmp snooping mrouter`
	`Vlan Ports`
	`-------- ------------`
	`1 Eth1/1(static)`

Related Commands	vlan <id> ip igmp snooping mrouter interface ethernet <id>
Notes	

17.10.5.21 show ip igmp snooping querier

	show ip igmp snooping querier [vlan <num>] Displays running IGMP snooping querier configuration on the VLANs.	
Syntax Description	vlan <num>	Displays the IGMP snooping querier configuration running on the specified VLAN
Default	N/A	
Configuration Mode	Any command mode	

History	3.3.4200	
	3.6.2100	Updated example
Example	switch (config) # show ip igmp snooping querier vlan 1 Snooping querier information for VLAN 1 IGMP Querier Present Querier IP address: 10.10.10.10 Query interval: 125 Response interval: 100 Group membership interval: 1 Robustness: 2 Version: 3	
Related Commands	vlan <id> ip igmp snooping querier	
Notes		

17.10.5.22 show ip igmp snooping querier-guard

	show ip igmp snooping querier-guard [interface {ethernet <port> | port-channel <lag-id> | mlag-port-channel <mlag-id>}] Shows status of IGMP query-guard mode and statistics of the denied IGMP query packets.	
Syntax Description	port	Ethernet port
	lag-id	LAG ID
	mlag-id	MLAG ID
Default	N/A	
Configuration Mode	config	
History	3.8.2000	
Example	switch (config) # show ip igmp snooping querier-guard Eth1/1: Querier Guard Mode : Enabled Denied IGMP Query Messages: 0 r-qa-sw-eth-86 [standalone: master] (config) #	
Related Commands	ip igmp snooping querier-guard interface <type> <id> ip igmp snooping querier-guard	
Notes		

17.10.5.23 show ip igmp snooping querier counters

	show ip igmp snooping querier counters [vlan <num> [group <group-id>]] Displays IGMP snooping querier counters.	
Syntax Description	vlan	Displays IGMP snooping querier counters on specific VLAN
	group	Multicast group IP address
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002	

Example	switch (config) # show ip igmp snooping querier counters vlan 10 Snooping querier counters for VLAN 10 General queries received: 0 General queries transmitted: 0 Group specific queries received : 0 Group specific queries transmitted : 0 Group source specific queries received : 0 Group source specific queries transmitted : 0 Leave messages received : 0 Leave messages transmitted : 0 V1/V2 reports received : 0 V1/V2 reports transmitted : 0 V3 reports received: 0 V3 reports transmitted: 0
Related Commands	
Notes	

17.10.5.24 show ip igmp snooping statistics

	show ip igmp snooping statistics Displays IGMP snooping statistical counters.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.1400	
	3.6.1002	Updated example
	3.6.2100	Updated example
Example	switch (config) # show ip igmp snooping statistics Snooping Statistics for VLAN 3770 General queries received : 3 General queries transmitted: 0 Group specific queries received : 0 Group specific queries transmitted: 0 Group and source specific queries received : 0 Group and source specific queries transmitted: 0 V1/V2 reports received : 0 V1/V2 reports transmitted : 0 Leave messages received : 0 Leave messages transmitted: 0 V3 reports received : 12 V3 reports transmitted : 0 Active Groups count: 2 Dropped packets: 0 Joins: 0	
Related Commands		
Notes		

17.10.5.25 show ip igmp snooping vlan

	show ip igmp snooping vlan {<vlan/vlan-range> | all} Displays IGMP configuration per VLAN or VLAN range.	
Syntax Description	vlan/vlan range	Displays IGMP VLAN configuration per specific VLAN or VLAN range
	all	Display IGMP VLAN configuration on all VLAN

Default	N/A
Configuration Mode	Any command mode
History	3.1.1400

Example	`switch (config) # show ip igmp vlan 1`
	`Vlan 1 configuration parameters:`
	`  IGMP snooping is enabled`
	`  IGMP version is V2`
	`  Snooping switch is acting as Non-Querier`
	`  mrouter static port list: Eth1/1`
	`  mrouter dynamic port list: none`

Related Commands	
Notes	

17.10.5.26 show ip igmp snooping profile

	show ip igmp snooping profile <profile_name> Show content of the specified IGMP profile	
Syntax Description	profile_name	Specified profile name.
Default	N/A	
Configuration Mode	Any command mode	
History	3.9.2100	

Example	`switch (config) # show ip igmp snooping profile proflie_1`
	
	`IGMP snooping profile proflie_1:`
	`  Count: 5`
	
	`  Configuration:`
	`  seq 10 permit 224.1.1.0/24 192.168.1.1/32`
	`  seq 20 deny 224.1.1.0/24 192.168.1.1/32`
	`  seq 53 permit 224.2.0.1/32 192.168.53.0/24`
	`  seq 54 permit 224.3.0.0/16 192.168.54.1/32`
	`  seq 65535 deny 0.0.0.0/0 0.0.0.0/0`	

Related Commands	ip igmp snooping profile	
Notes		

17.10.5.27 show ip igmp snooping filter

	show ip igmp snooping filter {interface {ethernet <port>[-<port>] | port-channel <lag-id>[-<lag-id>] | mlag-port-channel <mlag-id>[-<mlag-id>]} | vlan <num>[-<range>] } [detail[value]] | [statistics] Show statistics of the IGMP snooping filter.	
Syntax Description	port	Ethernet port
	lag-id	LAG ID
	mlag-id	MLAG ID
	vlan <num>[-<range>]	Specified VLAN or specified VLAN-range
	detail	IGMP filter detail information

	value	Specified number of Denied requested groups. Range: 10-100 Default: 10
	statistics	IGMP filter statistics
Default	N/A	
Configuration Mode	Any command mode	
History	3.9.2100	

Example	<pre>switch (config) # show ip igmp snooping filter interface ethernet 1/5 detail
	
	Eth1/5 IGMP Filters:
	  Status : Enabled
	  Profile Name: permitSpec
	
	  Profile statistics details:
	  VLAN range: 50
	
	  VLAN 50:
	  Denied requested groups : 1
	  Denied membership report packets : 1
	  Partially denied V3 membership report packets : 0
	
	  Denied group address list (last 10 entries):
	  239.1.2.22, 0.0.0.0
	
	  Max IGMP dynamic groups: 2
	
	  Max groups statistics details:
	  Denied requested groups : 1
	  Denied membership report packets : 1
	  Partially accepted V3 membership report packets: 0
	
	  Denied group address list (last 10 entries):
	
	  239.1.12.24, 0.0.0.0
	
	  Active groups: 1</pre>	

Related Commands		
Notes	• For IGMP Snooping filter feature to show denied group address list, only 50MB memory in total is allowed to be allocated. If 80% of 50MB is reached, the user will be notified. If 100% of 50MB is reached, the user will be notified and no more memory will be allowed to be allocated. Use “clear ip igmp snooping filter counters” command in the CLI to clear the memory. • For whole group record filtering, either by profile filtering or max-group limit, the following format content would be logged: **REJECT** IGMP report of (Source, Group) = (1st source, 239.1.12.33) from **Host** x.x.x.x due to **max-groups limit**. **REJECT** IGMP report of (Source, Group) = (source1, source2,... 239.1.12.33) from **Host** x.x.x.x due to **profile** (profile_name_xxx) filtering For profile filtering, it could be partial source address matching, i.e., some source addresses are filtered, while the group record remains not filtered. The related log is like the following format: “**REJECT these source address list** (src2, src4, ...) **out of IGMP report of Group** (239.239.0.18), from Host x.x.x.x due to profile (profile_name_xxx) filtering” • For partial source address matching, if there are some source address filtered from some group records of the report packet, then “Partially denied V3 membership report packets” will be updated accordingly	

18 Appendixes

The document contains the following appendixes:

- Appendix: Ethernet Storage Fabric (ESF)
- Appendix: Enhancing System Security According to NIST SP 800-131A
- Appendix: Feature Support per IC and CPU Type
- Appendix: Splunk Integration with NVIDIA Products
- Appendix: Show Commands Not Supported By JSON API
- Appendix: What Just Happened (WJH) Events

18.1 Appendix: Ethernet Storage Fabric (ESF)

Ethernet Storage Fabric (ESF) delivers performance and efficiency for scale-out storage and hyperconverged infrastructures. It leverages the speed, flexibility, and cost efficiency of Ethernet to provide the foundation for the fastest and most efficient storage networking infrastructure.

ESF runs on purpose-built switches which are optimized to deliver the highest levels of performance, lowest latency and zero packet loss, with unique form factor and storage aware features. Other capabilities of ESF include simultaneous handling of compute and storage traffic, future proofed with support for the NVMe over fabric protocol, support for file, block, and object storage, and it is best suited for scale-out storage and Hyperconverged infrastructures.

This section describes NVIDIA Ethernet Storage Fabric solution, its use cases, implementation and monitoring and debugging capabilities.

The most common deployment of ESF is a single rack of 6-18 servers, or in the case of HCI 6-18 appliances. The servers/appliances are connected in high availability architecture, utilizing MLAG, to two ToR SN2100/SN2010 half ``19 width Spectrum switches, enabling high availability in a single rack unit.

We will start with the setup/topology overview, followed by its Bill of Material and connectivity guidelines.

The following sections will describe the various ESF deployment manners available for the user:

1. CLI based configuration done one-by-one on all switches
2. Automation based configuration using Ansible
3. Using NEO as the management system

18.1.1 ESF Configuration using Ansible

Ansible is the de-facto standard for automation in the data center to enable efficiency, errorless mode of work and bottom line reaching lower TCO and faster TTM for deployments at scale
Detailed information on Ansible and the additional automation tools integrated with NVIDIA Onyx, please refer to the Automation chapter in this User Manual.
Here you can find a detailed guideline on Ansible deployment on top of the discussed topology— please refer to solution 1 described in this guide.
In this deployment guide we use a server/VM running Ansible, connecting the switches through the management network and configuring them using Ansible playbook composed of the Ansible modules building blocks available on NVIDIA Onyx page on Ansible.com.

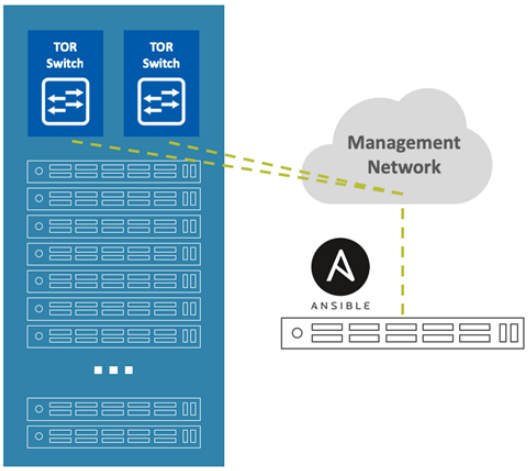

18.1.2 ESF Configuration Using CLI

Before starting the configuration process, make sure both switches have the same software version installed. To check the software version, run the "show version" command in the CLI.

It is recommended to upgrade both switches to the latest software release.

18.1.2.1 Switch Configuration

Run the following commands on both switches:

1. Enable LACP (required for the IPL):

```
sx01 (config) # lacp
```

2. Turn off spanning tree using this command (only if using version v3.6.6102 or earlier)

```
sx01 (config) # no spanning-tree
```

3. Enable IP routing:

```
sx01 (config) # ip routing
```

4. Enable MLAG protocol:

```
sx01 (config) # protocol mlag
```

5. Enable QoS globally:

```
sx01 (config) # dcb priority-flow-control enable force
```

18.1.2.2 IPL Configuration

Control traffic for the MLAG is sent over the IPL ports via a L3 interface (interface VLAN).
For high availability, it is recommended to have more than one physical link serving as the IPL, therefore the IPL is configured over LAG (port-channel).
It is recommended to use a VLAN ID that is not used within the subnet (4000 in this example) to avoid mixing the host traffic with the control traffic on this interface.
All VLANs are open on the IPL port. There is no need to configure this port, once an interface is mapped as "IPL", all VLANs are open on this port.
In this example, ports 1/35 and 1/36 are used for the IPL connectivity between the switches.
The IPL link may pass traffic upon MLAG port failures, but not under normal circumstances when all ports are in UP state.

Run the following commands on both switches:

```
sx01 (config) # interface port-channel 1
sx01 (config interface port-channel 1 ) # exit
sx01 (config) # interface ethernet 1/35 channel-group 1 mode active
sx01 (config) # interface ethernet 1/36 channel-group 1 mode active
sx01 (config) # vlan 4000
sx01 (config vlan 4000) # exit
sx01 (config) # interface vlan 4000
sx01 (config interface vlan 4000 ) # exit
sx01 (config) # interface port-channel 1 ipl 1
sx01 (config) # interface port-channel 1 dcb priority-flow-control mode on force
```

18.1.2.2.1 Configure IP address for the IPL link on both switches

1. Configure the following on one switch (e.g. sx01):

```
sx01 (config) # interface vlan 4000
sx01 (config interface vlan 4000) # ip address 10.10.10.1 255.255.255.0
sx01 (config interface vlan 4000) # ipl 1 peer-address 10.10.10.2
```

2. Configure the following on the second switch (e.g. sx02):

```
sx02 (config) # interface vlan 4000
sx02 (config interface vlan 4000) # ip address 10.10.10.2 255.255.255.0
sx02 (config interface vlan 4000) # ipl 1 peer-address 10.10.10.1
```

The IPL IP address should not be part of the management network, it could be any IP address and subnet that is not in use in the network. This address is not advertised outside the switch

18.1.2.3 MAGP Configuration

As stated in the previous chapter, MAGP configuration is required on the Spine switches when the fabric is utilizing L2 routing in the whole fabric. You can find more details about MAGP in the MAGP section of the UM.

To configure MAGP on the switches, you need to take the following steps on all spine switches used in your setup. In our use case we have one rack with two such switches:

18.1.2.3.1 Switch 1 Configuration

1. Create a VLAN interface.

```
switch (config)# interface vlan 20
switch (config interface vlan 20)#
```

2. Set an IP address to the VLAN interface.

```
switch (config interface vlan 20)# ip address 11.11.11.11 /8
```

3. Enable MAGP protocol globally.

```
switch (config)# protocol magp
```

18.1.2.3.2 Switch 2 Configuration

1. Create a VLAN interface:

```
switch (config)# interface vlan 20
switch (config interface vlan 20)#
```

2. Set an IP address to the VLAN interface.

```
switch (config interface vlan 20)# ip address 11.11.11.22 /8
```

3. Enable MAGP protocol globally.

```
switch (config)# protocol magp
```

4. Next steps (9-11) should be taken per VLAN (done for VLAN 10 below):Create a virtual router group for an IP interface.

```
switch (config interface vlan 20)# magp 10
```

5. Set a virtual router primary IP address.

```
switch (config interface vlan 20 magp 10)# ip virtual-router address 11.11.11.254
```

6. Set a virtual router primary MAC address.

```
switch (config interface vlan 20 magp 10)# ip virtual-router mac-address aa:bb:cc:dd:ee:ff
```

Verify the MAGP configuration.

```
switch (config)# show magp 10
```

The output in our setup will return the following:

```
MAGP 10
Interface vlan: 20
Admin state: Master
State: Enabled
Virtual IP: 11.11.11.254
Virtual MAC: aa:bb:cc:dd:ee:ff
```

18.1.2.4 MLAG Interface Configuration

MLAG configuration is very similar to port-channel configuration. It is recommended to keep the same port in each switch within the same mlag-port-channel (not a must). In this example, there are two MLAG ports, one for each host (host s1 is connected to mlag-port-channel 1 and host s2 is connected to mlag-port-channel 2).

The "mlag-port-channel" number is globally significant and must be the same on both switches.

1. Configure the following on both switches:

```
sx01 (config) # interface mlag-port-channel 1-2
sx01 (config interface port-channel 1-2 ) # exit
```

2. Set the mode (LACP or static) - Only one option is applicable:
   - To set the MLAG interface in static mode run:

```
sx01 (config) # interface ethernet 1/1 mlag-channel-group 1 mode on
sx01 (config) # interface ethernet 1/2 mlag-channel-group 2 mode on
```

   - To set the MLAG interface in LACP mode, run:

```
sx01 (config) #
interface ethernet 1/1 mlag-channel-group 1 mode active
sx01 (config) # interface ethernet 1/2 mlag-channel-group 2 mode active
```

   LACP mode 4 should be configured on the host side. Configuring LACP is similar in LAG and MLAG ports. LACP notifications arrive via the control protocol and not via the port physical status. It will show the remote system-id and may encounter configuration errors. LACP is very valuable, especially in large scale configurations with multiple MLAGs, as it helps detect any mismatched configurations in terms of connectivity.

3. Enable the two interfaces:

```
sx01 (config) #interface mlag-port-channel 1-2 no shutdown
```

4. To change any MLAG port parameter (e.g. MTU), enter the MLAG interface configuration mode and perform the change:

```
sx01 (config) # interface mlag-port-channel 1-2
sx01 (config interface mlag-port-channel 1-2 ) # mtu 9216 force
```

   Some operations may require "force" or manual disabling of the link.
5. To change the LAG/MLAG port speed, all interfaces should be removed from LAG/MLAG while changing the speed in the member interface configuration mode.
   It is recommended to configure the ports speed before adding the ports as members to the LAG/MLAG port, as once the ports are members in a LAG/MLAG the speed cannot be modified without removing the port from the LAG/MLAG.
6. To verify MLAG configuration and status, run the following commands:

```
sx01 [my-mlag-vip-domain: master] (config) # show mlag
Admin status: Enabled
Operational status: Up
Reload-delay: 30 sec
Keepalive-interval: 1 sec
System-id: F4:52:14:11:E5:38
```

```
MLAG Ports Configuration Summary:
Configured: 2
Disabled: 0
Enabled: 2
MLAG Ports Status Summary:
Inactive: 0
Active-partial: 0
Active-full: 2
MLAG IPLs Summary:
ID Group Vlan Operational Local Peer
Port-Channel Interface State IP address IP address
-------------------------------------------------------------------------------
1 Po1 4000 Up 10.10.10.1 10.10.10.2
```

7. To verify MLAG domain status, run:

```
sx01 [my-mlag-vip-domain: master] (config) # show mlag-vip
MLAG VIP
========
MLAG group name: my-mlag-vip-domain
MLAG VIP address: 10.209.28.200/24
Active nodes: 2
Hostname VIP-State IP Address
-----------------------------------------------------
sx01 master 10.209.28.50
sx02 standby 10.209.28.51
```

8. To see MLAG interfaces summary, run:

```
sx01 [my-mlag-vip-domain: master] (config) # show interfaces mlag-port-channel summary
MLAG Port-Channel Flags: D-Down, U-Up
P-Partial UP, S - suspended by MLAG
Port Flags: D - Down, P - Up in port-channel (members)
S - Suspend in port-channel (members), I - Individual
Group
Port-Channel Type Local Ports Peer Ports
(D/U/P/S) (D/P/S/I) (D/P/S/I)
-------------------------------------------------------------------------------
1 Mpo1(U) Static Eth1/1(P) Eth1/1(P)
2 Mpo2(U) Static Eth1/2(P) Eth1/2(P)
```

18.1.2.5 MLAG VIP Configuration

MLAG VIP (Virtual IP) is important for retrieving peer information.

The management network is used for keep-alive messages between the switches.

The MLAG domain must be unique name for each MLAG domain. In case you have more than one pair of MLAG switches on the same network, each domain (consist of two switches) should be configured with different name.

The IP address should be within the subnet of the management interface (mgmt0).

1. Configure the following on both switches:

```
sx01 (config)# mlag-vip my-mlag-vip-domain ip 10.209.28.200 /24 force
```

2. Set a virtual system MAC. The System MAC is used to identify the far-end switch used for the LACP System ID. It should be unicastrange.

```
switch (config)# mlag system-mac 00:00:5e:00:01:5d
```

   In case of an upgrade the MAC address is auto-calculated. For new MLAG installation, it must be added as configuration.
   The MLAG system-mac needs to be identical between both switches.

3. Enable MLAG globally, run:

```
switch config) # no mlag shutdown
```

18.1.2.6 Server Configuration

There are various options to configure a bond on the servers but not all bond modes are applicable. The supported bonding modes are as follows:

- balance-rr: mode 0
- balance-xor: mode 2
- 802.3ad (LACP): mode 4 (starting from 3.4.0000 MLNX-OS release)

Modes 1,3,5,6 were designed to work without LAG configured on the switch side, which limits support for all other modes. Configuring LAG on the switch side will break the solution.

For bonding modes which require LAG on the switch, MLAG must be configured when using redundant switches.

For the bonding modes which don't use LAG on the switch, two independent switches or non MLAG ports on MLAG switches are enough.

Linux Bonding Mode	Mode Number	LAG on switch requirement	Availability on MLAG interface
balance-rr	0	Yes	Yes
active-backup	1	No	No
balance-xor	2	Yes	Yes
broadcast	3	No	No
802.3ad	4	Yes (with LACP)	Yes
balance-tlb	5	No	No
balance-alb	6	No	No

Please refer to the below links for detailed examples:

- Example for Linux.
- Example for Windows 2012 (or above) where LBFO is configured via the OS.

In older Windows versions it is configured via the NIC driver configuration.

18.1.3 ESF Maintenance, Monitoring and Troubleshooting

18.1.3.1 MLAG Upgrade Procedure

To upgrade the MLAG cluster, the standby switch should be upgraded first, then (after reboot with the upgraded software) the slave will rejoin the MLAG cluster.

After that, the master can be upgraded.

When the master reboots with the upgraded software, the other standby node (which is running) becomes the master. After the old master reboots, it joins the cluster and then the configuration is set.

For a more detailed description of NVIDIA Onyx upgrade procedure, please refer to the following posts:

- HowTo Upgrade MLNX-OS Software on NVIDIA switch systems
- HowTo Upgrade MLNX-OS Software on an MLAG Switch Pair

18.1.3.2 Monitoring and Troubleshooting

This section provides information and tools to monitor and debug the deployed fabric.

It is recommended to ensure that the below conditions are followed:

1. Both switches are part of the same management subnet (connected to the same switch or more but on the same subnet).
2. The management network is connected on mgmt0 port.
3. The mlag-port-channel number is identical in both switches (recommended but not obligatory).
4. The same switch version is installed on both switches.
5. The IPL link is in UP state. try to ping the other switch via the IPL ping.
6. Align the MLAG interface mode on both the server and the switch.
   For example, if you select LACP mode on the MLAG interface (active), mode 4 should be configured on the bond interface.

Below are failure scenarios followed by monitoring and debug instructions.

The following scenarios are discussed:

- IPL link Down
- 'Inactive Ports' and 'Active-Partial' Status on the “show mlag” command
- Management Port is Down but IPL port is UP
- MLAG Cluster issues
- IPL issues
- MLAG port issues

18.1.3.2.1 IPL link Down

The IPL link should be configured as port-channel with 2 or more ports, but in some scenarios both ports may be in “Down” state. In this case only the master switch will pass traffic.

If we run “show mlag” command when only one “mlag-port-channel” port is configured, we will get the following:

Master:

```
mti-mar-sx04 [my-new-domain: master] (config) # show mlag
Admin status: Enabled
Operational status: Up
Reload-delay: 30 sec
Keepalive-interval: 1 sec
Upgrade-timeout: 60 min
System-mac: 00:00:5e:00:01:5d
MLAG Ports Configuration Summary:
Configured: 1
Disabled: 0
Enabled: 1
MLAG Ports Status Summary:
Inactive: 0
Active-partial: 0
Active-full: 1
MLAG IPLs Summary:
ID Group Vlan Operational Local Peer
Port-Channel Interface State IP address IP address
-------------------------------------------------------------------------------
1 Po1 4000 Up 10.10.10.2 10.10.10.1
MLAG Members Summary:
```

```
System-id State Hostname
-------------------------------------
E4:1D:2D:37:50:88 Up <mti-mar-sx04>
E4:1D:2D:37:54:88 Up mti-mar-sx03
mti-mar-sx04 [my-new-domain: master] (config) #
```

Standby:

```
mti-mar-sx03 [my-new-domain: standby] (config) # show mlag
Admin status: Enabled
Operational status: Up
Reload-delay: 30 sec
Keepalive-interval: 1 sec
Upgrade-timeout: 60 min
System-mac: 00:00:5e:00:01:5d
MLAG Ports Configuration Summary:
Configured: 1
Disabled: 0
Enabled: 1
MLAG Ports Status Summary:
Inactive: 0
Active-partial: 0
Active-full: 1
MLAG IPLs Summary:
ID Group Vlan Operational Local Peer
Port-Channel Interface State IP address IP address
-------------------------------------------------------------------------------
1 Po1 4000 Up 10.10.10.1 10.10.10.2
MLAG Members Summary:
System-id State Hostname
-------------------------------------
E4:1D:2D:37:54:88 Up <mti-mar-sx03>
E4:1D:2D:37:50:88 Up mti-mar-sx04
mti-mar-sx03 [my-new-domain: standby] (config) #
```

When shutting down the IPL port on the master switch:

```
mti-mar-sx04 [my-new-domain: master] (config) # interface port-channel 1 shutdown
mti-mar-sx04 [my-new-domain: master] (config) # show mlag
Admin status: Enabled
Operational status: Up
Reload-delay: 30 sec
Keepalive-interval: 1 sec
Upgrade-timeout: 60 min
System-mac: 00:00:5e:00:01:5d
MLAG Ports Configuration Summary:
Configured: 1
Disabled: 0
Enabled: 1
MLAG Ports Status Summary:
Inactive: 0
Active-partial: 0
Active-full: 1
MLAG IPLs Summary:
ID Group Vlan Operational Local Peer
Port-Channel Interface State IP address IP address
-------------------------------------------------------------------------------
1 Po1 4000 Down 10.10.10.2 10.10.10.1
MLAG Members Summary:
System-id State Hostname
-------------------------------------
E4:1D:2D:37:50:88 Up <mti-mar-sx04>
E4:1D:2D:37:54:88 Down mti-mar-sx03
mti-mar-sx04 [my-new-domain: master] (config) #
```

Standby switch:

```
mti-mar-sx03 [my-new-domain: standby] (config) # show mlag
Admin status: Enabled
Operational status: Down
Reload-delay: 30 sec
Keepalive-interval: 1 sec
Upgrade-timeout: 60 min
System-mac: 00:00:5e:00:01:5d
MLAG Ports Configuration Summary:
Configured: 1
Disabled: 1
Enabled: 0
```

```
MLAG Ports Status Summary:
Inactive: 0
Active-partial: 0
Active-full: 1
MLAG IPLs Summary:
ID Group Vlan Operational Local Peer
Port-Channel Interface State IP address IP address
-------------------------------------------------------------------------------
1 Po1 4000 Down 10.10.10.1 10.10.10.2
MLAG Members Summary:
System-id State Hostname
-----------------------------------
E4:1D:2D:37:54:88 Peering <mti-mar-sx03>
E4:1D:2D:37:50:88 Down mti-mar-sx04
mti-mar-sx03 [my-new-domain: standby] (config) #
```

18.1.3.2.2 'Inactive Ports' and 'Active-Partial' Status on the “show mlag” command

By default, all ethernet ports are admin UP, while the mlag-port-channels are down, as in most cases the full network configuration is done first and then the mlag-port-channel is enabled. Make sure to enable the ports when creating mlag-port-channel and adding ethernet interface to it (either static or LACP).

Note: When one port is down, it doesn't mean that the whole mlag-port-channel is down.

MLAG Ports Status Summary:

- Inactive - all ports in the mlag-port-channel are down (on both switches).
- Active-partial - some ports are down (example below, on one switch)
- Active-full - normal condition, all is good.

When one mlag-port-channel is down, we will see the following output:

```
mti-mar-sx03 [my-new-domain: master] (config) # interface mlag-port-channel 10 shutdown
mti-mar-sx03 [my-new-domain: master] (config) # show mlag
Admin status: Enabled
Operational status: Up
Reload-delay: 30 sec
Keepalive-interval: 1 sec
Upgrade-timeout: 60 min
System-mac: 00:00:5e:00:01:5d
MLAG Ports Configuration Summary:Configured: 1
Disabled: 0
Enabled: 1
MLAG Ports Status Summary:Inactive: 0
Active-partial: 1
Active-full: 0
MLAG IPLs Summary:
ID Group Vlan Operational Local Peer
Port-Channel Interface State IP address IP address
-------------------------------------------------------------------------------
1 Po1 4000 Up 10.10.10.1 10.10.10.2
MLAG Members Summary:
System-id State Hostname
-----------------------------------
E4:1D:2D:37:54:88 Up <mti-mar-sx03>E4:1D:2D:37:50:88 Up mti-mar-sx04
mti-mar-sx03 [my-new-domain: master] (config) #
```

To enable it:

```
mti-mar-sx03 [my-new-domain: master] (config) # interface mlag-port-channel 10 no shutdown
mti-mar-sx03 [my-new-domain: master] (config) # show mlag
Admin status: Enabled
Operational status: Up
Reload-delay: 30 sec
Keepalive-interval: 1 sec
Upgrade-timeout: 60 min
System-mac: 00:00:5e:00:01:5d
MLAG Ports Configuration Summary:
Configured: 1
Disabled: 0
Enabled: 1
MLAG Ports Status Summary:
Inactive: 0
Active-partial: 0
Active-full: 1
MLAG IPLs Summary:
ID Group Vlan Operational Local Peer
Port-Channel Interface State IP address IP address
```

```
--------------------------------------------------------------------------------
1 Po1 4000 Up 10.10.10.1 10.10.10.2
MLAG Members Summary:
System-id State Hostname
-----------------------------------
E4:1D:2D:37:54:88 Up <mti-mar-sx03>
E4:1D:2D:37:50:88 Up mti-mar-sx04
mti-mar-sx03 [my-new-domain: master] (config) #
```

18.1.3.2.2.1 Management Port is Down but IPL port is UP

When there is no ping between the two servers on mgmt0 (e.g. mgmt0 port is Down, or any management switch problem that blocks traffic between the switches on mgmt0) - both switches will pass traffic.

There is no mentioning of the second switch in the cluster.

The “show mlag” and “show mlag-vip” output will look like this:

```
mti-mar-sx04 [my-new-domain: master] (config) # show mlag
Admin status: Enabled
Operational status: Up
Reload-delay: 30 sec
Keepalive-interval: 1 sec
Upgrade-timeout: 60 min
System-mac: 00:00:5e:00:01:5d
MLAG Ports Configuration Summary:
Configured: 1
Disabled: 0
Enabled: 1
MLAG Ports Status Summary:
Inactive: 0
Active-partial: 0
Active-full: 1
MLAG IPLs Summary:
ID Group Vlan Operational Local Peer
Port-Channel Interface State IP address IP address
--------------------------------------------------------------------------------
1 Po1 4000 Up 10.10.10.2 10.10.10.1
MLAG Members Summary:
System-id State Hostname
-----------------------------------
E4:1D:2D:37:50:88 Up <mti-mar-sx04>
E4:1D:2D:37:54:88 Up -
mti-mar-sx04 [my-new-domain: master] (config) #
mti-mar-sx04 [my-new-domain: master] (config) # show mlag-vip
MLAG VIP
========
MLAG group name: my-new-domain
MLAG VIP address: 10.20.2.205/24
Active nodes: 1
Hostname VIP-State IP Address
-----------------------------------------------------
mti-mar-sx04 master 10.20.2.54
mti-mar-sx04 [my-new-domain: master] (config) #
```

18.1.3.2.2.2 MLAG Cluster Issues

After adding the two switches to the cluster, wait for a few seconds. One switch will become Master, while the other one will become the slave. When performing remove/add/cluster change operations, always wait for the switch to go to “standalone master” before continuing.

Run "show mlag-vip"

```
mti-mar-sx03 [my-mlag-vip-domain: master] (config) # show mlag-vip
MLAG VIP
========
MLAG group name: my-mlag-vip-domain
MLAG VIP address: 10.20.2.205/24
Active nodes: 2
Hostname VIP-State IP Address
-----------------------------------------------------
mti-mar-sx03 master 10.20.2.53
mti-mar-sx04 standby 10.20.2.54
mti-mar-sx03 [my-new-domain: master] (config) #
```

Verify that the two switches are in the cluster. The other MLAG switch must reflect the same information.

If one switch does not see this MLAG-Domain do the following:

Run "show ip route":

```
mti-mar-sx03 [my-mlag-vip-domain: master] (config) # show ip route
VRF Name: default
---------------------------
Destination Mask Gateway Interface Source Distance/Metric
default 0.0.0.0 10.20.0.251 mgmt0 DHCP 0/0
10.20.0.0 255.255.0.0 0.0.0.0 mgmt0 direct 0/0
10.10.10.0 255.255.255.0 0.0.0.0
```

The management subnet must only point out of the MGMT port. inband management is acceptable. If there is a conflict, the MGMT Keep alive is sent out on the wrong port and not advertised to another switch.

In case the switch still does not see the cluster: The MGMT keep alive is broadcast to a well known multicast DNS group – 224.0.0.251. Check to see if both switches are advertising to this group. It is likely that the mgmt. port will see a lot of traffic. This output will need to be captured and analyzed.

```
mti-mar-sx03 [my-mlag-vip-domain: master] (config) # tcpdump -i mgmt0
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on mgmt0, link-type EN10MB (Ethernet), capture size 96 bytes
06:42:15.330780 IP mti-mar-sx03.mti.labs.mlnx.mdns > 224.0.0.251.mdns: 0 [2a] PTR (Cache flush)? _tcn_MLAG-
DOMAIN._tcp.local. (117)
```

This is a transmission from master to the multicast group. Before we have a master, both switches will see this frame, and both will transmit it. After the cluster is formed, only the master will transmit this. If this frame is not seen, the cluster will not form.

18.1.3.2.3 IPL issues

IPL Link needs to be up for MLAG peer ports and sync data to be available. The IPL VLAN is local to the MLAG switches and can be any number. VLAN 4000 or higher is typically used for control vlans and is recommended.

The “show mlag” command shows IPL link state and other valuable information.

The IPL link needs to be Up. Both switches must be in Up State in the “Member” summary. Peering or down are not a good state. Peering could be a transient state but should move to UP eventually.

```
mti-mar-sx03 [my-mlag-vip-domain: master] (config) # show mlag
Admin status: Enabled
Operational status: Up
Reload-delay: 30 sec
Keepalive-interval: 1 sec
Upgrade-timeout: 60 min
System-mac: 00:00:5e:00:01:5d << Both switches should show the same System MAC Address
MLAG Ports Configuration Summary:
Configured: 1
Disabled: 0
Enabled: 1
MLAG Ports Status Summary:
Inactive: 0
Active-partial: 0
Active-full: 1
MLAG IPLs Summary:
ID Group Vlan Operational Local Peer
Port-Channel Interface State IP address IP address
-------------------------------------------------------------------------------
1 Po1 4000 Up 10.10.10.1 10.10.10.2
MLAG Members Summary:
System-id State Hostname
-----------------------------------
E4:1D:2D:37:54:88 Up <mti-mar-sx03>
E4:1D:2D:37:50:88 Up mti-mar-sx04
```

In case IPL is up and still member ports are not visible, try ping the remote IPL interface. Ping the local switch and then the MLAG Peer switch IPL IP address. If ping doesn't go through use tcpdump to debug this case. In case link is up and ping is lossy, check for traffic on the IPL interface. During normal operation, IPL traffic is a few frames per second at the most. If you see a lot of traffic, it is likely an indication of a loop in the setup.

```
switch (config) # tcpdump -i vlan4000
```

The other usual suspects are checking if both sides are set to static, or LACP. Check interface transceiver for matching serial numbers to identify cabling issues.

18.1.3.2.4 MLAG Port Issues

A healthy MLAG should show all ports as UP (P) and MLAG must be (U).

```
mti-mar-sx03 [my-mlag-vip-domain: master] (config) # show interface mlag-port-channel summary
MLAG Port-Channel Flags: D-Down, U-Up
P-Partial UP, S - suspended by MLAG
Port Flags: D - Down, P - Up in port-channel (members)
S - Suspend in port-channel (members), I - Individual
Group
Port-Channel Type Local Ports Peer Ports
(D/U/P/S) (D/P/S/I) (D/P/S/I)
------------------------------------------------------------------------------
1 Mpo1(U) LACP Eth1/10(P) Eth1/10(P)
mti-mar-sx03 [my-mlag-vip-domain: master] (config) #
```

"Partial" means that all ports are down on the MLAG-peer switch side. This could be a result of interface MLAG being shut on the remote side or mlag protocol shut on remote side.

```
mti-mar-sx03 [my-mlag-vip-domain: master] (config) # show interface mlag-port-channel summary
MLAG Port-Channel Flags: D-Down, U-Up
P-Partial UP, S - suspended by MLAG
Port Flags: D - Down, P - Up in port-channel (members)
S - Suspend in port-channel (members), I - Individual
Group
Port-Channel Type Local Ports Peer Ports
(D/U/P/S) (D/P/S/I) (D/P/S/I)
------------------------------------------------------------------------------
1 Mpo1(P) LACP Eth1/10(P) Eth1/10(D)
```

Peer ports not being visible means that ports in the MLAG-Peer switch are either not added in the MLAG or there are cluster issues.

```
mti-mar-sx03 [my-mlag-vip-domain: master] (config) # show interface mlag-port-channel summary
MLAG Port-Channel Flags: D-Down, U-Up
P-Partial UP, S - suspended by MLAG
Port Flags: D - Down, P - Up in port-channel (members)
S - Suspend in port-channel (members), I - Individual
Group
Port-Channel Type Local Ports Peer Ports
(D/U/P/S) (D/P/S/I) (D/P/S/I)
------------------------------------------------------------------------------
1 Mpo1(P) LACP Eth1/10(P)
SX1012-B [MLAG-DOMAIN: master] (config) #
```

If the physical port shows (S) that could result from either receiving no PDUs from the remote side or by receiving a PDU that doesn't match what is being received on other members of the MLAG port-channel

Check the LACP counters to see continuous increment of counters, both sent and receive must increment. One every second for fast retransmit and one every 30 seconds for slow retransmit.

```
mti-mar-sx03 [my-mlag-vip-domain: master] (config) # show lacp counters
```

```
LACPDUs Marker Marker Response LACPDUs
Port Sent Recv Sent Recv Sent Recv Illegal Unknown
-------------------------------------------------------------------------------
...
Mlag-port-channel: 1
------------------
1/10 0 0 0 0 35 27 0 0
```

In case the lacp counters are incrementing and port is still down, then check the SID received on different port of the MLAG. They should match across all MLAG ports.

```
mti-mar-sx03 [my-mlag-vip-domain: master] (config) #show lacp interfaces neighbors
Flags:
A - Device is in Active mode
P - Device is in Passive mode
MLAG channel group 1 neighbors
Port 1/10
----------
Partner System ID : e4:1d:2d:37:48:80 (This is the System-ID received on this port from the remote switch. It must
match for all ports connected to the same switch)
Partner System priority : 32768
Flags : A
LACP Partner Port Priority : 32768
LACP Partner Oper Key : 13845 (LACP OPER KEY must match across all ports in the same MLAG port-channel)
LACP Partner Port State : 0xbc
Port State Flags Decode
-----------------------
Activity : Active
Aggregation State : Aggregation, Sync, Collecting, Distributing,
```

To check the SID used by the NVIDIA switch use this command:

```
mti-mar-sx03 [my-mlag-vip-domain: master] (config) # show lacp interfaces mlag-port-channel 1 system-identifier
Priority: 32768
MAC: 00:00:5e:00:01:06
```

Check the lacp property across all ports in an MLAG:

```
mti-mar-sx03 [my-mlag-vip-domain: master] (config) # show lacp interfaces eth 1/10
Port : 1/10
-------------
Port State = Bundle
MLAG Channel Group : 1
Pseudo mlag-port-channel = Mpo1
LACP port-priority = 32768
LACP Rate = Slow
LACP Activity : Active
LACP Timeout : Short
Aggregation State : Aggregation, Sync, Collecting, Distributing,
LACP Port Admin Oper Port Port
Port State Priority Key Key Number State
-------------------------------------------------------------------
1/7 Bundle 32768 29001 29001 0x7 0x0
(This is what we advertise to the remote switch- the Admin and Oper keys must match across all ports in a port-
channel)
```

18.1.4 ESF Setup Examples

18.1.4.1 Single Rack with Two Switches Connected in MLAG

In this setup, we cover the most common deployment scenario and most cost-effective solution: Two switches in a single rack configured with MLAG, providing high availability for the connected servers (as described in the below diagrams).

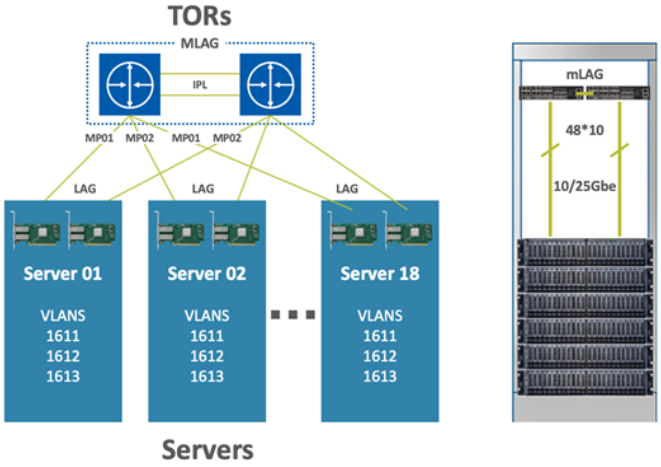

To leverage the high availability and connectivity to the L3 cloud, Multi-Active Gateway Protocol (MAGP) is used, resolving the default gateway problem when a host is connected to a set of switch routers (SRs) via MLAG with no LACP control (MAGP is a NVIDIA proprietary protocol that implements active-active VRRP). The network functionality in that case requires that each SR is an active default gateway router to the host, thus reducing hops between the SRs and directly forwarding IP traffic to the L3 cloud regardless which SR traffic comes through.

In ESF deployment in a single rack, the ToR switches' router ports are configured for connectivity with the external network.

To get a detailed overview of the MLAG terminology and its architecture, please refer to the MLAG section in this user manual.

18.1.4.1.1 Bill of Materials

As described in the diagram above (two switches in a Rack running MLAG) the fabric in this solution is built with the following components:

Component	Quantity	Description
Leaf Switch	2	SN2010 Spectrum based 25GbE/100GbE, 1U Open Ethernet Switch with Onyx, 18 SFP28 and 4 QSFP28 ports, 2 Power Supplies (AC), short depth, x86 quad core, P2C airflow, Rail Kit must be purchased separately, RoHS6
Servers	Max 18	N/A
Uplinks	2	N/A
Network Adapters	2 per server	ConnectX-5 Dual-Port SFP28 Port, PCIe 3.0 x16, tall bracket, ROHS R6
Leaf-Server Cable	1 per server	SFP28 25GbE Passive Copper Cable
Leaf-Leaf Cable (IPL)	2 per rack	QSFP28 100GbE Passive Copper Cable

18.1.4.1.2 Physical Network Connectivity

The setup connectivity configuration will be as follows:

- 2 NVIDIA Spectrum SN2010 (used as the TOR switches)
- 2 X 100GbE uplink ports for the WAN/LAN connectivity (Up to 18 nodes in a rack and a total of 4 x 100GbE uplink ports)
- 2 X 100GbE ports (on each switch) for switch connectivity (IPL) using 2 X QSFP28 100GbE Passive Copper Cables
- Dedicated management port on each switch connected to the Switch Management Network
- Single 25GbE connection from the server to each TOR switch by using the SFP28 25GbE Passive Copper Cable

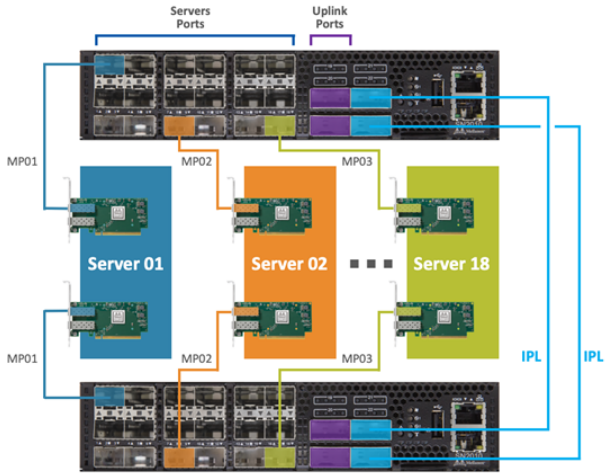

18.1.4.2 Scale-out Common Deployments

When moving from a single rack deployment into a Leaf-Spine deployment where the ToR switches of each rack are connected to spine switches, there are two major deployment options:

1. Whole fabric L2 with MLAG configured on the ToR and spine switches, and the Spine switches deploy MAGP.
2. L2 up to the ToR switches and L3 routing between the ToR and spine switches.

Please refer to the following community post for BGP deployment on top of MLAG in a leaf-spine topology.

18.2 Appendix: Enhancing System Security According to NIST SP 800-131A

Our switch systems, by default, work with NIST SP 800-131A, as described in the table below.

This appendix describes how to enhance the security of a system in order to comply with the NIST SP 800-131A standard. This standard is a document which defines cryptographically “acceptable” technologies. This document explains how to protect against possible cryptographic vulnerabilities

in the system by using secure methods. Because of compatibility issues, this security state is not the default of the system and it should be manually set.

> Some protocols, however, cannot be operated in a manner that complies with the NIST SP 800-131A standard.

Component	Configuration	Command
HTTP	HTTP disabled	no web http enable
HTTPS	HTTPS enabled	no web https enable
	SSL ciphers = TLS1.2	web https ssl ciphers all
	SSL renegotiation disabled	web https ssl renegotiation enable
SSH	SSH version = 2	ssh server min-version 1
	SSH ciphers = aes256-ctr, aes192-ctr, aes128-ctr, aes128-gcm@openssh.com, aes256-gcm@openssh.com	no ssh server security strict

18.2.1 Web Certificate

The OS supports signature generation of sha256WithRSAEncryption, sha1WithRSAEncryption self-signed certificates, and importing certificates as text in PEM format.

To configure a default certificate:

1. Create a new sha256 certificate.

```
switch (config) # crypto certificate name <cert name> generate self-signed hash-algorithm sha256
```

> For more details and parameters refer to the command “crypto certificate name”.

2. Show crypto certificate detail.

```
switch (config) # show crypto certificate detail
```

Search for “signature algorithm” in the output.

3. Set this certificate as the default certificate. Run:

```
switch (config) # crypto certificate default-cert name <cert name>
```

To configure default parameters and create a new certificate:

1. Define the default hash algorithm.

```
switch (config) # crypto certificate generation default hash-algorithm sha256
```

2. Generate a new certificate with default values.

```
switch (config) # crypto certificate name <cert name> generate self-signed
```

> When no options are selected, the generated certificate uses the default values for each field.

To test strict mode connect to the WebUI using HTTPS and get the certificate. Search for “signature algorithm”.

> There are other ways to configure the certificate to sha256. For example, it is possible to use “certificate generation default hash-algorithm” and then regenerate the certificate using these default values.

> It is recommended to delete browsing data and previous certificates before retrying to connect to the WebUI.

> Make sure not to confuse “signature algorithm” with “Thumbprint algorithm”.

18.2.2 SNMP

SNMPv3 supports configuring username, authentication keys and privacy keys. For authentication keys it is possible to use MD5 or SHA. For privacy keys AES or DES are to be used.

To configure strict mode, create a new user with HMAC-SHA1-96 and AES-128. Run:

```
switch (config) # snmp-server user <username> v3 auth sha <password1> priv aes-128 <password2>
```

To verify the user in the CLI, run:

```
switch (config) # show snmp user
```

> To test strict mode, configure users and check them using the CLI, then run an SNMP request with the new users.

> SNMPv1 and SNMPv2 are not considered to be secure. To run in strict mode, only use SNMPv3.

18.2.3 HTTPS

By default, the OS supports HTTPS encryption using TLS1.2 only. Working in TLS1.2 mode also bans MD5 ciphers which are not allowed per NIST 800-131a. In strict mode, the switch supports encryption with TLS1.2 only with the following supported ciphers:

- RSA_WITH_AES_128_CBC_SHA256
- RSA_WITH_AES_256_CBC_SHA256
- DHE_RSA_WITH_AES_128_CBC_SHA256
- DHE_RSA_WITH_AES_256_CBC_SHA256
- TLS_RSA_WITH_AES_128_GCM_SHA256
- TLS_RSA_WITH_AES_256_GCM_SHA384
- TLS_DHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_DHE_RSA_WITH_AES_256_GCM_SHA384

To enable all encryption methods, run:

```
switch (config) # web https ssl ciphers all
```

To enable only TLS ciphers (enabled by default), run:

```
switch (config) # web https ssl ciphers TLS
```

To enable HTTPS strict mode, run:

```
switch (config) # web https ssl ciphers TLS1.2
```

To verify which encryption methods are used, run:

```
switch (config)# show web
Web User Interface:
 Web interface enabled: yes
 HTTP enabled: yes
 HTTP port: 80
 HTTP redirect to HTTPS: no
 HTTPS enabled: yes
 HTTPS port: 443
 HTTPS ssl-ciphers: TLS1.2
 HTTPS certificate name: default-cert
 Listen enabled: yes
 No Listen Interfaces.

 Inactivity timeout: disabled
 Session timeout: 2 hr 30 min
 Session renewal: 30 min

Web file transfer proxy:
 Proxy enabled: no

Web file transfer certificate authority:
 HTTPS server cert verify: yes
 HTTPS supplemental CA list: default-ca-list
```

On top of enabling HTTPS, to prevent security breaches HTTP must be disabled.

To disable HTTP, run:

```
switch (config) # no web http enable
```

18.2.4 Code Signing

Code signing is used to verify that the data in the image is not modified by any third-party. The operating system supports signing the image files with SHA256, RSA2048 using GnuPG.

> Strict mode is operational by default.

18.2.5 SSH

The SSH server on the switch by default uses secure ciphers only, message authentication code (MAC), key exchange methods, and public key algorithm. When configuring SSH server to strict mode, the aforementioned security methods only use approved algorithms as detailed in the NIST 800-181A specification and the user can connect to the switch via SSH in strict mode only.

To enable strict security mode, run the following:

```
switch (config) # ssh server security strict
```

> The following ciphers are disabled for SSH when strict security is enabled:
> - 3des-cbc
> - aes256-cbc
> - aes192-cbc
> - aes128-cbc
> - rijndael-cbc@lysator.liu.se

The no form of the command disables strict security mode.

Make sure to configure the SSH server to work with minimum version 2 since 1 is vulnerable to security breaches.

To configure min-version to strict mode, run:

```
switch (config) # ssh server min-version 2
```

> Once this is done, the user cannot revert back to minimum version 1.

18.2.6 LDAP

By default, the switches support LDAP encryption SSL version 3 or TLS1.0 up to TLS1.2. The only banned algorithm is MD5 which is not allowed per NIST 800-131a. In strict mode, the switch supports encryption with TLS1.2 only with the following supported ciphers:

- DHE-DSS-AES128-SHA256

- DHE-RSA-AES128-SHA256
- DHE-DSS-AES128-GCM-SHA256
- DHE-RSA-AES128-GCM-SHA256
- DHE-DSS-AES256-SHA256
- DHE-RSA-AES256-SHA256
- DHE-DSS-AES256-GCM-SHA384
- DHE-RSA-AES256-GCM-SHA384
- ECDH-ECDSA-AES128-SHA256
- ECDH-RSA-AES128-SHA256
- ECDH-ECDSA-AES128-GCM-SHA256
- ECDH-RSA-AES128-GCM-SHA256
- ECDH-ECDSA-AES256-SHA384
- ECDH-RSA-AES256-SHA384
- ECDH-ECDSA-AES256-GCM-SHA384
- ECDH-RSA-AES256-GCM-SHA384
- ECDHE-ECDSA-AES128-SHA256
- ECDHE-RSA-AES128-SHA256
- ECDHE-ECDSA-AES128-GCM-SHA256
- ECDHE-RSA-AES128-GCM-SHA256
- ECDHE-ECDSA-AES256-SHA384
- ECDHE-RSA-AES256-SHA384
- ECDHE-ECDSA-AES256-GCM-SHA384
- ECDHE-RSA-AES256-GCM-SHA384
- AES128-SHA256
- AES128-GCM-SHA256
- AES256-SHA256
- AES256-GCM-SHA384

To enable LDAP strict mode, run the following:

```
switch (config) # ldap ssl mode {start-tls | ssl}
```

Both modes operate using SSL. The different lies in the connection initialization and the port used.

18.3 Appendix: Feature Support per IC and CPU Type

The following table lists which features are supported by which IC family and CPU type.

New features added on release 3.6.81xx and beyond are supported on Spectrum-based switches only.

Feature	NVIDIA Spectrum Family
Image Maintenance via ONIE	Yes

Feature	NVIDIA Spectrum Family
IPv6	Yes
JSON	Yes
OpenFlow 1.0	Yes
OpenFlow 1.3	Yes
PIM	Yes
PTP	Yes
QoS RED & ECN	Yes
S&F config	Yes
Signal Degradation Monitoring	Yes
Shared Buffers	Yes
Storm Control	Yes
Telemetry (histograms and threshold)	Yes
User Defined Keys	Yes
VXLAN	Yes

18.4 Appendix: Splunk Integration with NVIDIA Products

Splunk automatically clusters millions of log records in real time back into their patterns and finds connections between those patterns to form the baseline flows of each software individually, thus enables you to search, monitor and analyze that data to discover powerful insights across multiple use cases.

This appendix provides a guide on the first steps with Splunk and helps you to begin enjoying reduced time in detecting and resolving production problems.

18.4.1 Getting Started with Splunk

1. Download Splunk and extract the Splunk Enterprise version. (Splunk software is available as an RPM or TGZ.)

2. Create a Splunk User /group. Run:

```
[root@server] groupadd splunk
[root@server] useradd -d /opt/splunk -m -g splunk splunk
```

3. Splunk installation. Run:

```
[root@server] tar -xzvf splunk-7.0.0-c8a78efdd40f-Linux-x86_64.tgz
[root@server] ls
```

4. A new folder called Splunk is created.

```
[root@server] cp -rp splunk/* /opt/splunk/
[root@server] chown -R splunk: /opt/splunk/
[root@server] su - splunk
[splunk@server] cd bin
[splunk@server] ./splunk start --accept-license
```

Now you can access your Splunk WebUI at http://IP:8000/ or http://hostname:8000/. You need to make sure that port 8000 is open in your server firewall.

18.4.2 Switch Configuration

In this example we are not using the default UDP port 514 to show that any other port can be also used.

5. In order to add a task, the switch must be configured to send logs to our Splunk server. Run:

```
switch > enable
switch # configure terminal
switch (config) # show snmp
SNMP enabled:         yes
SNMP port:            161
System contact:
System location:

Read-only communities:
   public


Read-write communities:
   (none)

Interface listen enabled: yes
No Listen Interfaces.

switch (config) # snmp-server host 10.212.23.1 informs port 8597
switch (config) # snmp-server host 10.212.23.1 traps port 8597
switch (config) # snmp host 10.212.23.1 informs 8597
switch (config) # snmp host 10.212.23.1 traps  8597


Summary configuration:


switch (config) # show running-config
## Logging configuration
##
   logging 10.212.23.1
   logging 10.212.23.1 port 8597
   logging 10.212.23.1 trap info
   logging 10.212.23.1 trap override class events priority err
   logging monitor events notice
   logging receive
## SNMP configuration
no snmp-server host 10.209.21.221 disable
snmp-server host 10.209.21.221 traps port 8597 version 2c
no snmp-server host 10.212.23.1 disable
snmp-server host 10.212.23.1 traps port 8597 version 2c 8597
```

18.4.3 Adding a Task

6. The first screen encountered after signing into the Splunk WebUI includes the “Add Data” icon.

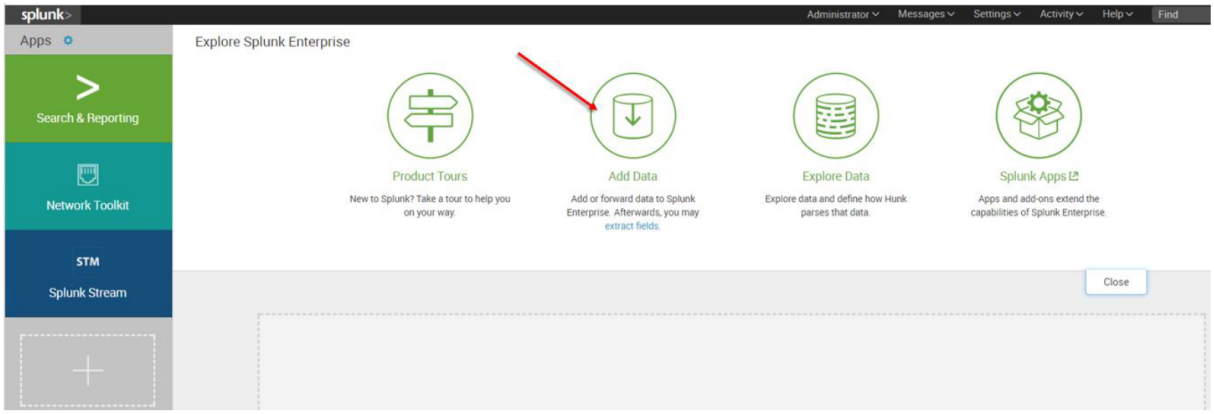

7. The “Add Data” tab opens up with three options: Upload, Monitor, and Forward. Here our task is to monitor a folder, so we click Monitor. to proceed

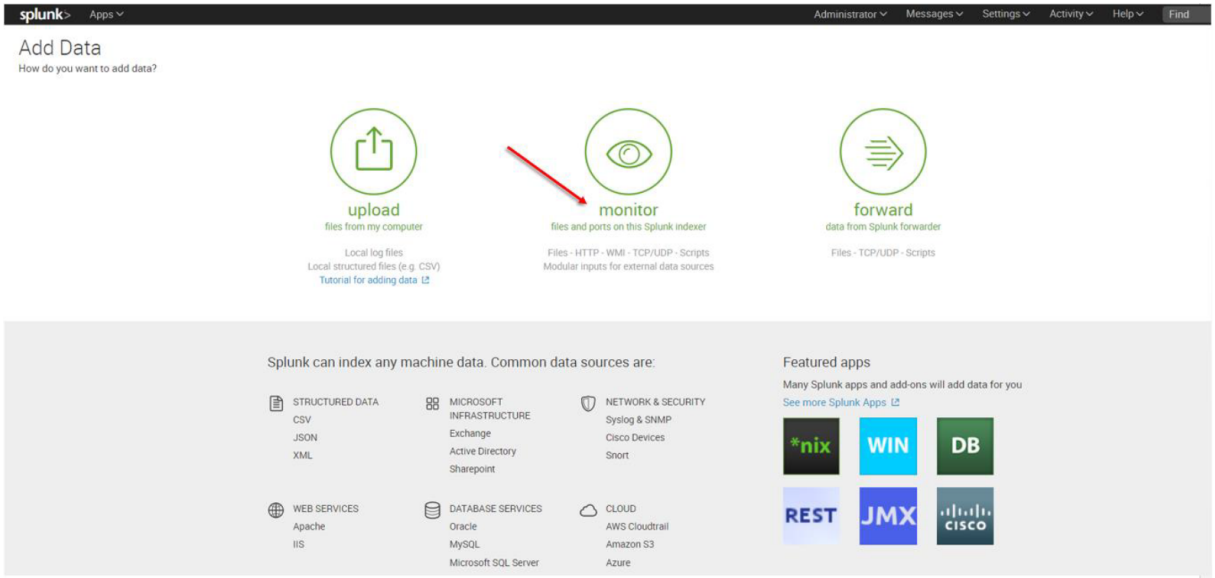

In the Monitor option, the following four categories are available:

- File & Directories – monitor files/folders
- HTTP Event Collector – monitor data streams over HTTP
- TCP/UDP – monitor service ports
- Scripts – monitor scripts

18.4.4 Retrieving Data from TCP and UDP Ports

8. Per our current purpose, we choose TCP/UDP option.

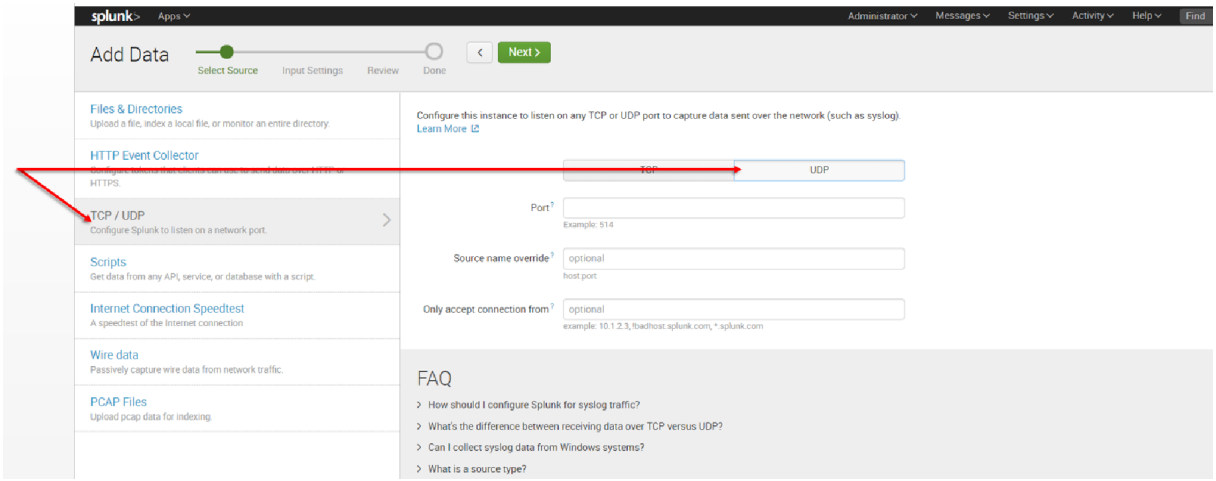

9. Click the TCP or UDP button to choose between a TCP or UDP input, and enter a port number in the “Port” field.

10. In the “Source name override” field, enter a new source name to override the default source value, if required.

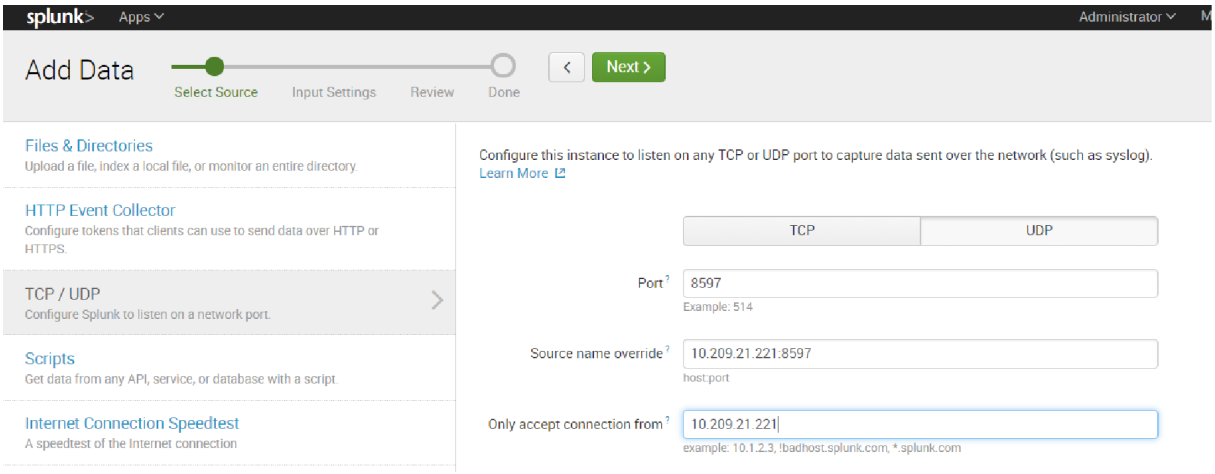

11. Click “Next” to continue to the Input Settings page where we will create a new source type called Mellanox-Switch.

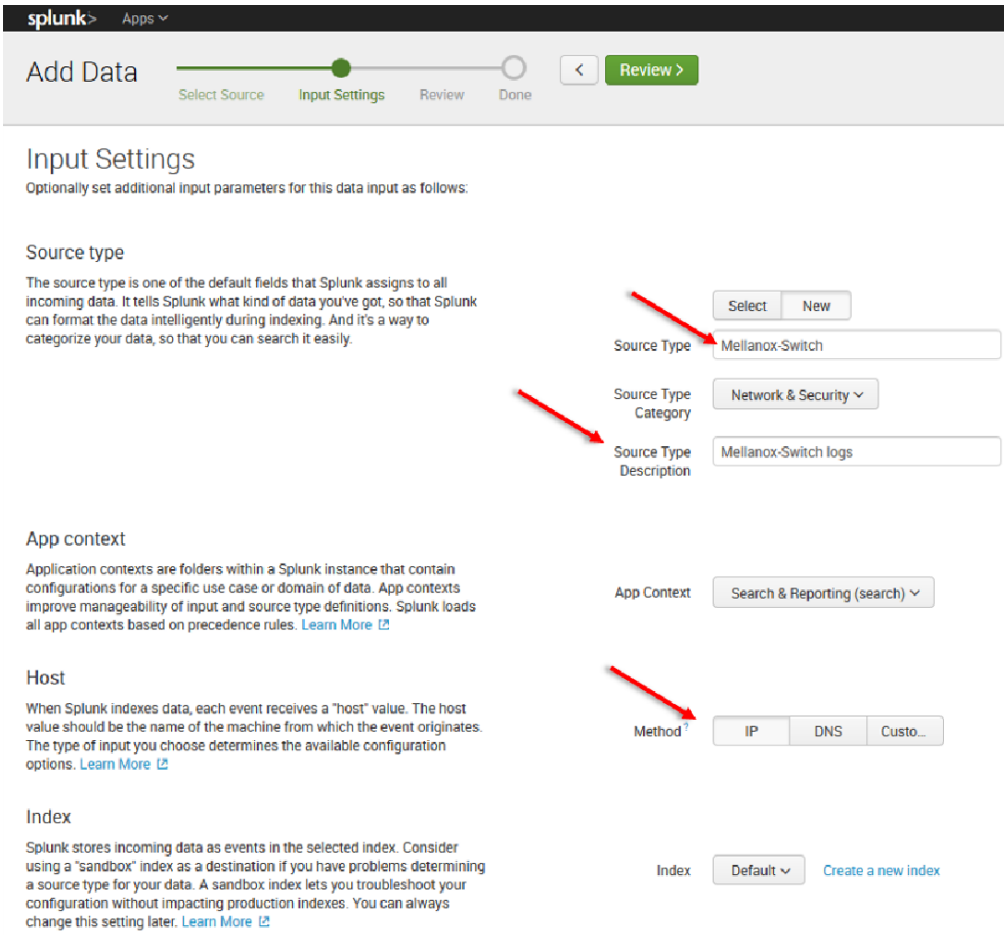

12. Click Next > Review > Done > Start Searching

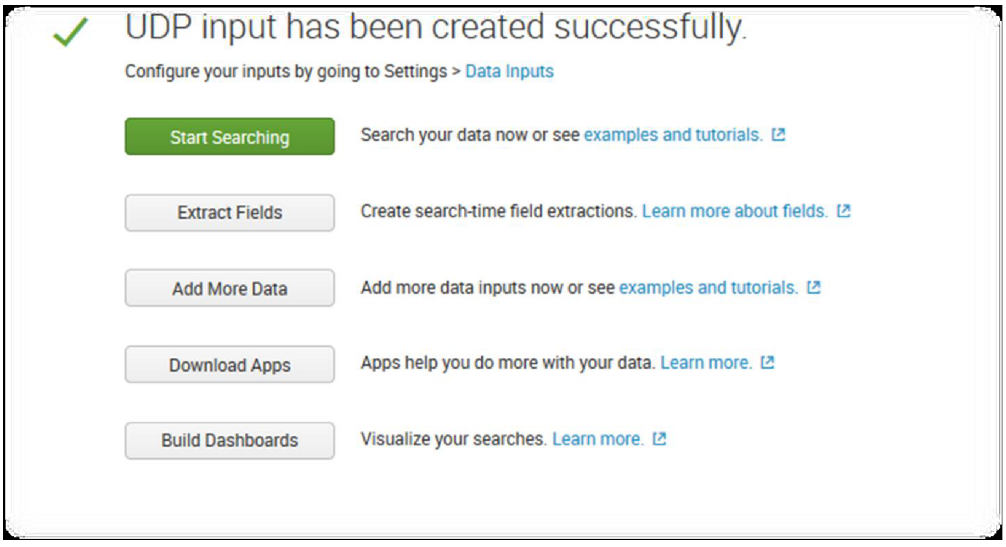

18.4.5 SNMP Input to Poll Attribute Values and Catch Traps

SNMP represents an incredibly rich source of data that you can get into Splunk for visibility across a very diverse IT landscape.

SNMP agents may also send notifications, called Traps, to an SNMP trap listening daemon.

18.4.5.1 Getting Started

Browse to Splunkbase and download the SNMP Modular Input from https://splunkbase.splunk.com/app/1537/.

To install, simply untar the file to SPLUNK_HOME/etc/apps and restart Splunk.

18.4.5.2 Configuration

Login to the Splunk WebUI and go to Manager > Add Data > Monitor > SNMP > New, and set up your input data.

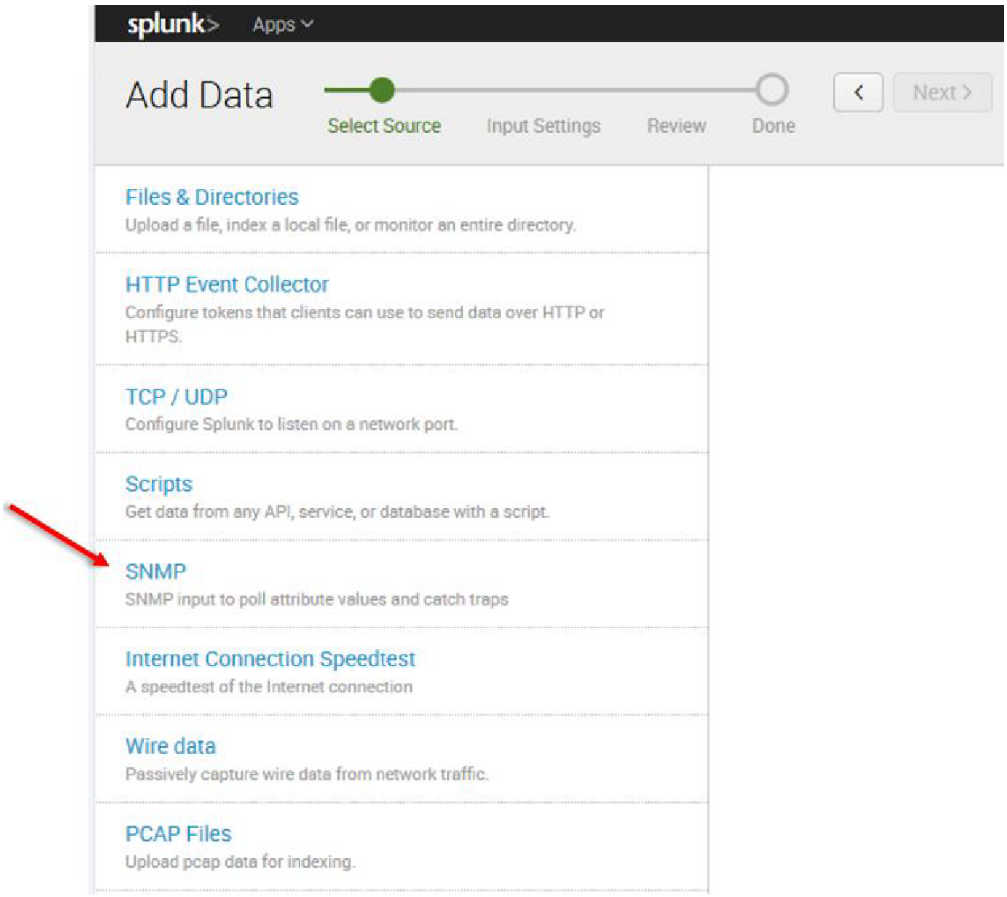

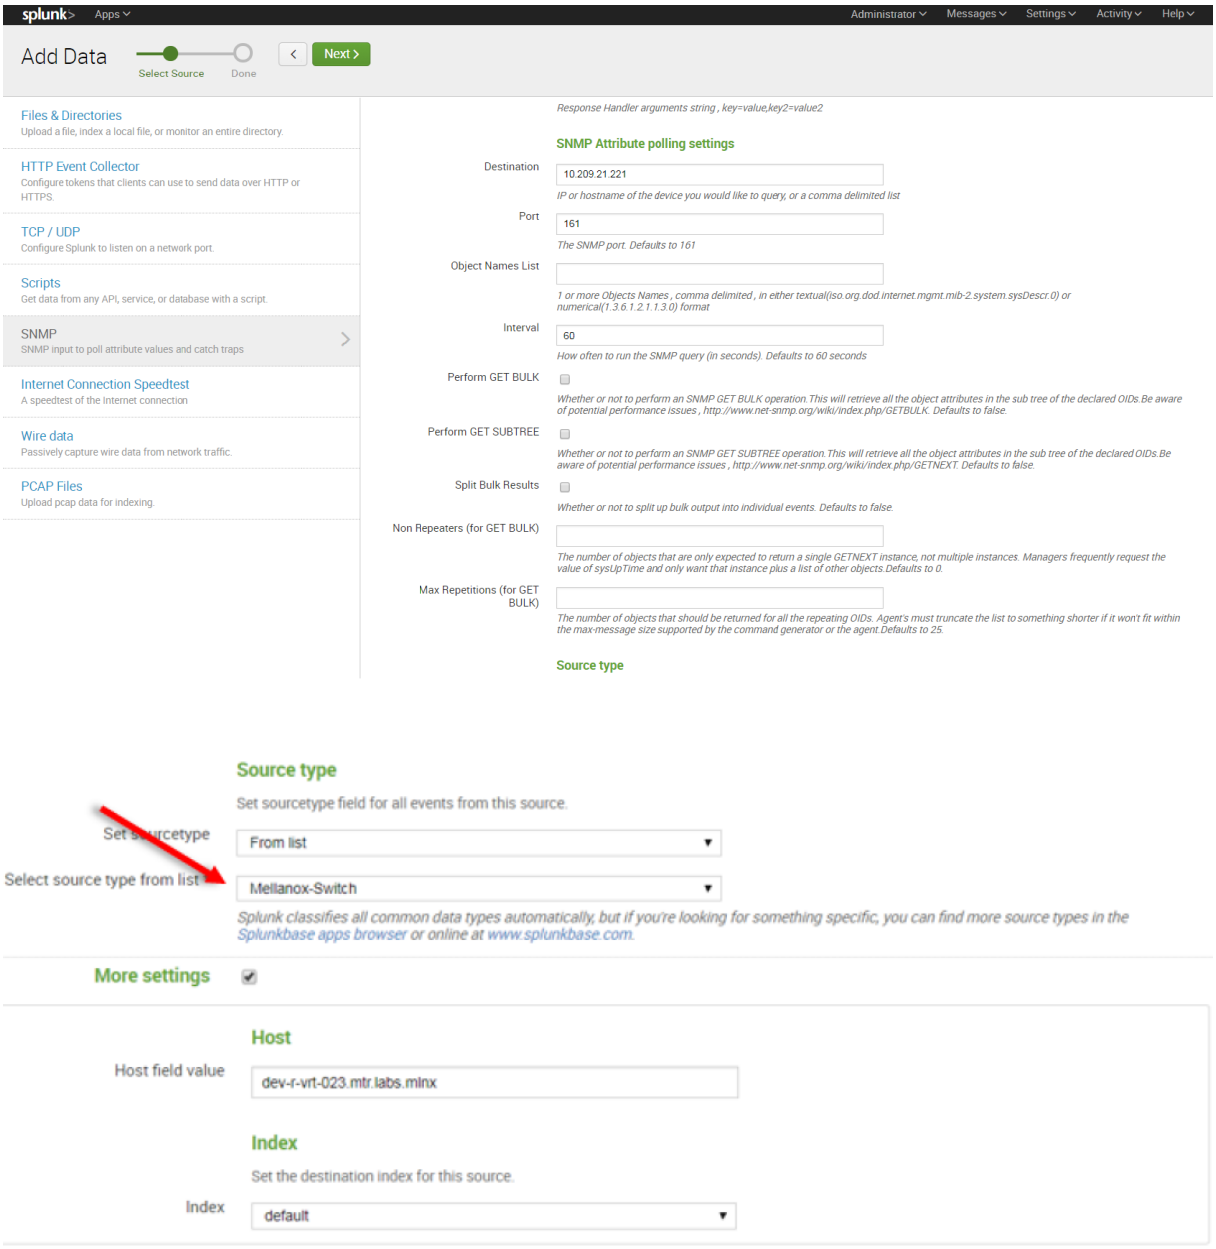

13. After configuration is complete it is recommend to run Mellanox-Switch again: Search > Data Summary > Sourcetypes > Mellanox-Switch.

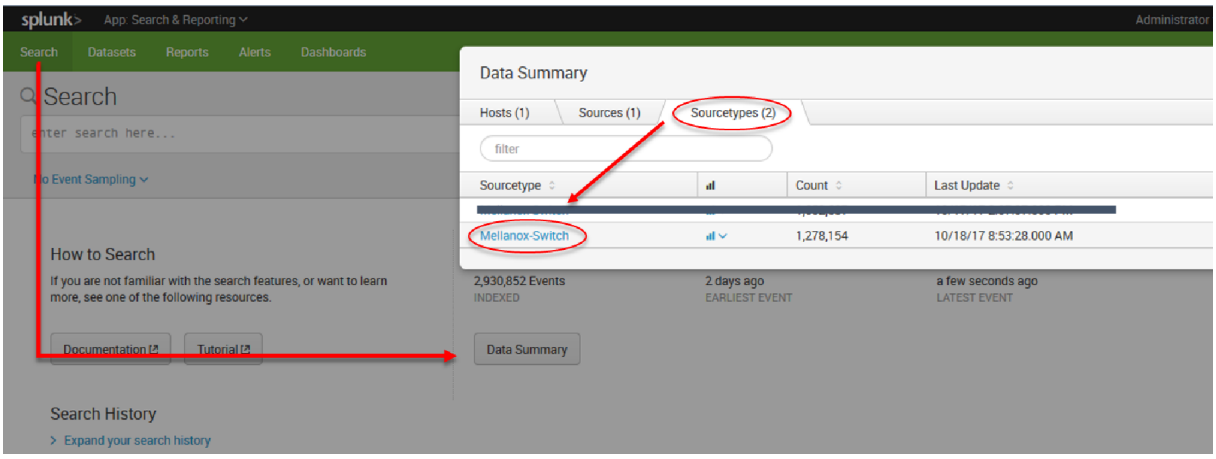

14. Select “Mellanox-Switch” and “Add to search”.

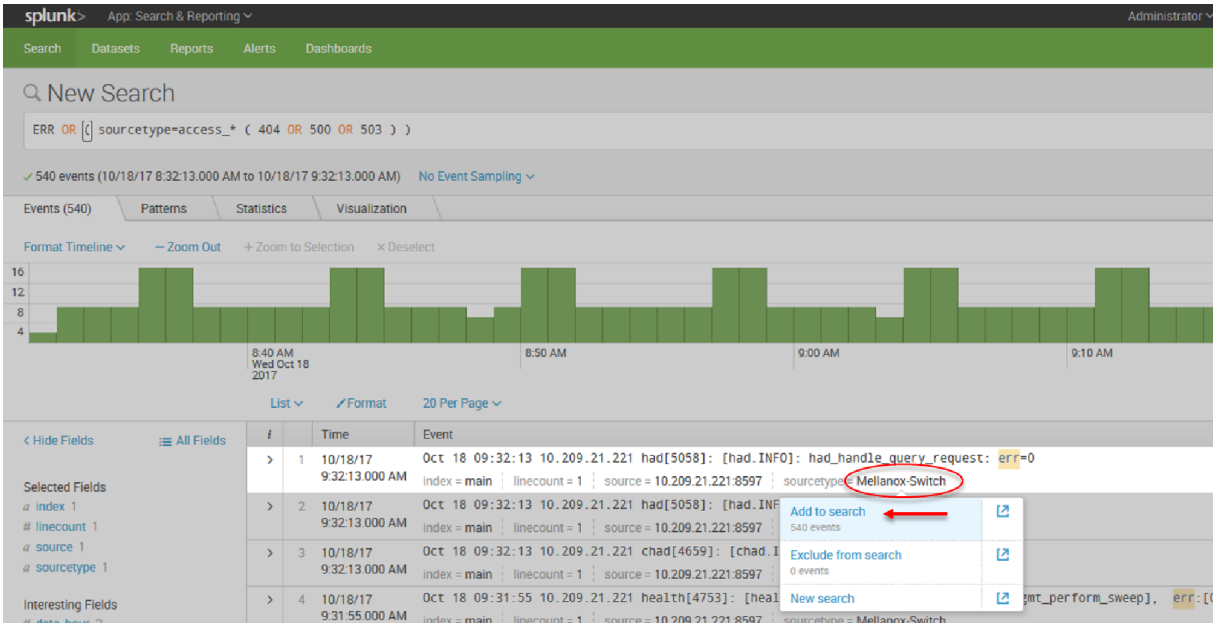

15. You can add to search any value that is relevant for you.

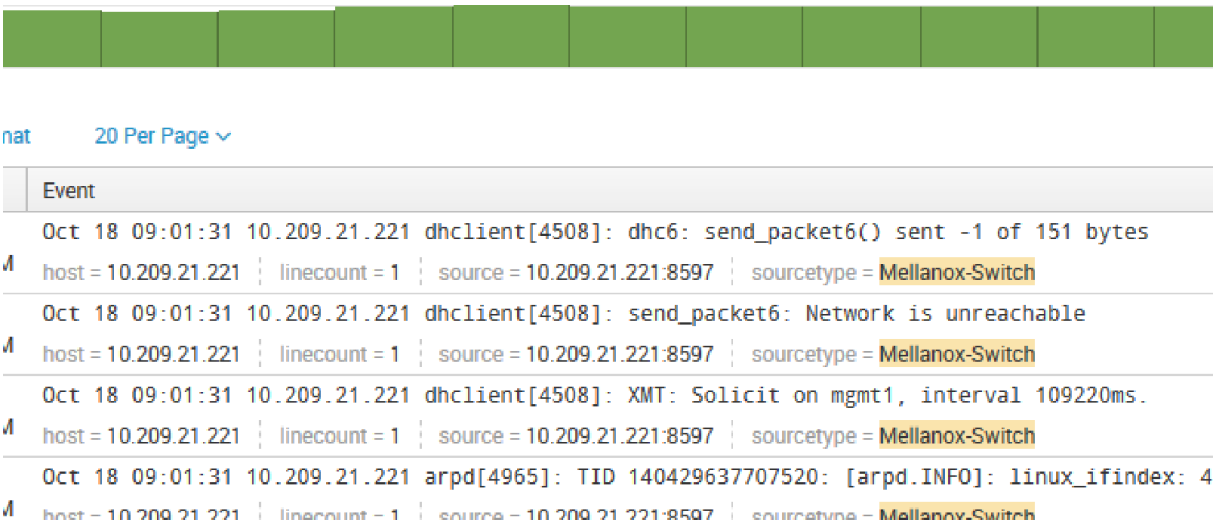

Patterns can be viewed not on real time and you can create alert on most repeatable events.

18.5 Appendix: Show Commands Not Supported By JSON API

Configuration Management
show configuration text files *
show files debug-dump *
show files stats *
Logging
show log
show log continuous
show log continuous matching *
show log continuous not matching *
show log debug
show log debug continuous
show log debug continuous matching *
show log debug continuous not matching *
show log debug files
show log debug files *
show log debug files * matching *
show log debug files * not matching *
show log debug matching *
show log debug not matching *
show log files
show log files *
show log files * matching *
show log files * not matching *
show log matching *
show log not matching *
Scheduled Jobs
show jobs
show jobs *
User Management and Security
show users history

show users history username *
User Interfaces
show cli
show cli max-sessions
show cli num-sessions
show terminal

18.6 Appendix: What Just Happened (WJH) Events

Drop Reason Group	Drop Reason	Comment
L1	Port admin down	Port Down Reason
L1	Auto-negotiation failure	Port Down Reason
L1	Logical mismatch with peer link	Port Down Reason
L1	Link training failure	Port Down Reason
L1	Peer is sending remote faults	Port Down Reason
L1	Bad signal integrity	Port Down Reason
L1	Cable/transceiver is not supported	Port Down Reason
L1	Cable/transceiver is unplugged	Port Down Reason
L1	Calibration failure	Port Down Reason
L1	Port state changes	Counter
L1	Symbol error	Counter
L1	CRC error	Counter
Forwarding	MLAG port isolation	Not supported for port isolation implemented with system ACL
Forwarding	Destination MAC is reserved (DMAC=01-80-C2-00-00-0x)	
Forwarding	VLAN tagging mismatch	
Forwarding	Ingress VLAN filtering	
Forwarding	Ingress spanning tree filter	
Forwarding	Unicast MAC table action discard	Currently not supported
Forwarding	Multicast egress port list is empty	
Forwarding	Port loopback filter	
Forwarding	Source MAC is multicast	
Forwarding	Source MAC equals destination MAC	
Forwarding	Non-routable packet	Currently not supported
Forwarding	Blackhole route	
Forwarding	Unresolved next-hop	
Forwarding	Blackhole ARP/neighbor	

Drop Reason Group	Drop Reason	Comment
Forwarding	IPv6 destination in multicast scope FFx0:/16	
Forwarding	IPv6 destination in multicast scope FFx1:/16	
Forwarding	Non-IP packet	
Forwarding	Unicast destination IP but non-unicast destination MAC	
Forwarding	Destination IP is loopback address	
Forwarding	Source IP is multicast	
Forwarding	Source IP is in class E	
Forwarding	Source IP is loopback address	
Forwarding	Source IP is unspecified	
Forwarding	Checksum or IP ver or IPv4 IHL too short	
Forwarding	Multicast MAC mismatch	
Forwarding	Source IP equals destination IP	
Forwarding	IPv4 source IP is limited broadcast	
Forwarding	IPv4 destination IP is local network (destination = 0.0.0.0/8)	
Forwarding	IPv4 destination IP is link local	
Forwarding	Ingress router interface is disabled	
Forwarding	Egress router interface is disabled	
Forwarding	IPv4 routing table (LPM) unicast miss	
Forwarding	IPv6 routing table (LPM) unicast miss	
Forwarding	Router interface loopback	
Forwarding	Packet size is larger than MTU	
Forwarding	TTL value is too small	
Forwarding	Overlay switch - source MAC is multicast	
Forwarding	Overlay switch - source MAC equals destination MAC	
Forwarding	Decapsulation error	
ACL	Ingress port ACL	
ACL	Ingress router ACL	
ACL	Egress port ACL	
ACL	Egress router ACL	
Buffer	Tail drop	
Buffer	WRED	

19 Document Revision History

Version	Date	Description
3.10.45xx LTS 3.10.44xx LTS	July 2024 January 2024	There are no changes to this version of the user manual. For further information on bug fixes and improvements, see the release notes of this software version.
3.10.43xx LTS	July 2023	Removed 'route-map-name' option from Networking command. Added: a note in MLAG Virtual System-MAC section.
3.10.43xx LTS	June 2023	There are no changes to this version of the user manual. For further information on bug fixes and improvements, see the release notes of this software version.
3.10.42xx LTS	February 2023	There are no changes to this version of the user manual. For further information on bug fixes and improvements, see the release notes of this software version.
3.10.41xx LTS	November 2022	Added note in the section Getting Started.Updated the section "Upgrading MLAG-STP Setup"
3.10.40xx	October 2022	Removed: • The command "ip l3" command • Puppet Agent section • Virtal Machine section • The command "show interfaces ethernet transceiver counters details
3.10.31xx	August 2022	Updated the module-type command.
3.10.30xx	July 2022	Added the ip filter reset-to-default-rules command.
		Updated: • The command "show interfaces ethernet transceiver counters" Removed: • show interfaces ethernet transceiver counters details
3.10.22xx	May 2022	There are no changes to this version of the user manual. For further information on bug fixes and improvements, see the release notes of this software version.
3.10.21xx	April 2022	There are no changes to this version of the user manual. For further information on bug fixes and improvements, see the release notes of this software version.

Revision	Date	Description
3.10.20xx	March 2022	Added: • The command "ldap nested-group-search" • The command "ldap nested-group-depth" • The command "ldap nested-group-count" • Note in the command "system secure-mode enable" • The command "show ptp timeout counters" • The command "clear ptp timeout counters" • The command "ptp monitor logging enable" • The command "no ptp monitor logging enable" • The command "show ptp monitor" • The command "show ptp monitor phc" • The command "ptp monitor interval" • The command "no ptp monitor interval" • The command "ptp monitor interval phc" • The command "no ptp monitor interval phc" Updated: • The command "show ldap"
3.10.12xx	January 2022	There are no changes to this version of the user manual. For further information on bug fixes and improvements, see the release notes of this software version.

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. Neither NVIDIA Corporation nor any of its direct or indirect subsidiaries and affiliates (collectively: “NVIDIA”) make any representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer (“Terms of Sale”). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer’s own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, “MATERIALS”) ARE BEING PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason

whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, and Mellanox are trademarks and/or registered trademarks of NVIDIA Corporation and/or Mellanox Technologies Ltd. in the U.S. and in other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright