



Mellanox MLNX-OS[®] User Manual for VPI

Rev 4.60

Software Version 3.6.3004



NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT (PRODUCT(S)) AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES AS-IS, WITH ALL FAULTS OF ANY KIND AND SOLELY FOR THE PURPOSE OF AIDING THE CUSTOMER IN TESTING APPLICATIONS THAT USE THE PRODUCTS IN DESIGNATED SOLUTIONS. THE CUSTOMER'S MANUFACTURING TEST ENVIRONMENT HAS NOT MET THE STANDARDS SET BY MELLANOX TECHNOLOGIES TO FULLY QUALIFY THE PRODUCT(S) AND/OR THE SYSTEM USING IT. THEREFORE, MELLANOX TECHNOLOGIES CANNOT AND DOES NOT GUARANTEE OR WARRANT THAT THE PRODUCTS WILL OPERATE WITH THE HIGHEST QUALITY. ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT ARE DISCLAIMED. IN NO EVENT SHALL MELLANOX BE LIABLE TO CUSTOMER OR ANY THIRD PARTIES FOR ANY DIRECT, INDIRECT, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING, BUT NOT LIMITED TO, PAYMENT FOR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY FROM THE USE OF THE PRODUCT(S) AND RELATED DOCUMENTATION EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



Mellanox Technologies
350 Oakmead Parkway Suite 100
Sunnyvale, CA 94085
U.S.A.
www.mellanox.com
Tel: (408) 970-3400
Fax: (408) 970-3403

© Copyright 2017. Mellanox Technologies Ltd All Rights Reserved

Mellanox®, Mellanox logo, Accelio®, BridgeX®, CloudX logo, CompustorX®, ConnectIB®, ConnectX®, CoolBox®, CORE-Direct®, EZchip®, EZchip logo, EZappliance®, EZdesign®, EZdriver®, EZsystem®, GPUDirect®, InfiniHost®, InfiniBridge®, InfiniScale®, Kotura®, Kotura logo, Mellanox CloudRack®, Mellanox CloudXMellanox®, Mellanox Federal Systems®, Mellanox HostDirect®, Mellanox Multi-Host®, Mellanox Open Ethernet®, Mellanox OpenCloud®, Mellanox OpenCloud Logo®, Mellanox PeerDirect®, Mellanox ScalableHPC®, Mellanox StorageX®, Mellanox TuneX®, Mellanox Connect Accelerate Outperform logo, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, NP-1c®, NP-2®, NP-3®, Open Ethernet logo, PhyX®, PlatformX®, PSIPHY®, SiPhy®, StoreX®, SwitchX®, Tilera®, Tilera logo, TestX®, TuneX®, The Generation of Open Ethernet logo, UFM®, Unbreakable Link®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies Ltd.

All other trademarks are property of their respective owners

For the most updated list of Mellanox trademarks visit <http://www.mellanox.com/page/trademarks>

Table of Contents

Table of Contents	3
List of Tables	13
List of Figures	16
Document Revision History	18
About this Manual	43
Chapter 1 Introduction	47
1.1 System Features	47
1.2 Ethernet Features	48
1.3 InfiniBand Features	49
1.4	49
Chapter 2 Getting Started	51
2.1 Configuring the Switch for the First Time	51
2.1.1 Re-Running the Wizard	58
2.2 Starting the Command Line (CLI)	58
2.3 Starting the Web User Interface (WebUI)	59
2.4 Licenses	63
2.4.1 Installing MLNX-OS® License (CLI)	63
2.4.2 Installing MLNX-OS License (Web)	64
2.4.3 Retrieving a Lost License Key	66
2.4.4 Commands	68
Chapter 3 User Interfaces	73
3.1 Command Line Interface Overview	73
3.1.1 CLI Modes	73
3.1.2 Syntax Conventions	74
3.1.3 Getting Help	74
3.1.4 Prompt and Response Conventions	75
3.1.5 Using the “no” Form	76
3.1.6 Parameter Key	77
3.1.7 CLI Pipeline Operator Commands	78
3.1.8 CLI Shortcuts	81
3.2 Web Interface Overview	82
3.2.1 Setup Menu	83
3.2.2 System Menu	84
3.2.3 Security Menu	85
3.2.4 Ports Menu	85

3.2.5	Status Menu	86
3.2.6	IB SM Mgmt	87
3.2.7	Fabric Inspector	87
3.2.8	ETH Mgmt	88
3.2.9	IP Route	89
3.3	Secure Shell (SSH)	89
3.3.1	Adding a Host and Providing an SSH Key	89
3.3.2	Retrieving Return Codes when Executing Remote Commands	90
3.4	Management Information Bases (MIBs)	90
3.5	Commands	94
3.5.1	CLI Session	94
3.5.2	Banner	106
3.5.3	SSH	114
3.5.4	Remote Login	132
3.5.5	Web Interface	135

Chapter 4 System Management **152**

4.1	Management Interface	152
4.1.1	Configuring Management Interfaces with Static IP Addresses	152
4.1.2	Configuring IPv6 Address on the Management Interface	152
4.1.3	Dynamic Host Configuration Protocol (DHCP)	153
4.1.4	Default Gateway	153
4.1.5	In-Band Management	153
4.1.6	Configuring Hostname via DHCP (DHCP Client Option 12)	154
4.1.7	Commands	155
4.2	NTP, Clock & Time Zones	206
4.2.1	NTP Authenticate	206
4.2.2	NTP Authentication Key	206
4.2.3	Commands	207
4.3	Unbreakable Links	225
4.3.1	Link Level Retransmission (LLR)	225
4.3.2	Configuring Phy Profile & LLR	226
4.3.3	Commands	228
4.4	Virtual Protocol Interconnect (VPI)	234
4.4.1	Commands	236
4.5	System Profile	239
4.5.1	Commands	240
4.6	Software Management	242
4.6.1	Upgrading MLNX-OS Software	242
4.6.2	Upgrading MLNX-OS Software on Director Switches	246
4.6.3	Upgrading MLNX-OS HA Groups	247
4.6.4	Deleting Unused Images	248

4.6.5	Downgrading MLNX-OS Software	248
4.6.6	Upgrading System Firmware	252
4.6.7	Image Maintenance via Mellanox ONIE	254
4.6.8	Commands	257
4.7	Configuration Management	270
4.7.1	Saving a Configuration File	270
4.7.2	Loading a Configuration File	270
4.7.3	Restoring Factory Default Configuration	271
4.7.4	Managing Configuration Files	271
4.7.5	Commands	274
4.8	Logging	306
4.8.1	Monitor	306
4.8.2	Remote Logging	306
4.8.3	Commands	307
4.9	Debugging	330
4.9.1	Commands	331
4.10	Event Notifications	351
4.10.1	Supported Events	351
4.10.2	SNMP Trap Notifications	352
4.10.3	Terminal Notifications	353
4.10.4	Email Notifications	353
4.10.5	Commands	355
4.11	Telemetry	374
4.11.1	Commands	374
4.12	mDNS	391
4.12.1	Commands	392
4.13	User Management and Security	393
4.13.1	User Accounts	393
4.13.2	Authentication, Authorization and Accounting (AAA)	393
4.13.3	System Secure Mode	395
4.13.4	Commands	397
4.14	Cryptographic (X.509, IPSec) and Encryption	445
4.14.1	System File Encryption	445
4.15	Scheduled Jobs	463
4.15.1	Commands	463
4.16	Statistics and Alarms	473
4.16.1	Commands	473
4.17	Chassis Management	494
4.17.1	System Health Monitor	494
4.17.2	Power Management	501
4.17.3	Monitoring Environmental Conditions	504

4.17.4 USB Access	506
4.17.5 Unit Identification LED	507
4.17.6 High Availability (HA)	507
4.17.7 System Reboot	511
4.17.8 Commands	513
4.18 Network Management Interfaces	547
4.18.1 SNMP	547
4.18.2 JSON API	558
4.18.3 XML API	563
4.18.4 Commands	564
4.19 Puppet Agent	589
4.19.1 Setting the Puppet Server	589
4.19.2 Accepting the Switch Request	589
4.19.3 Installing Modules on the Puppet Server	590
4.19.4 Writing Configuration Classes	590
4.19.5 Supported Configuration Capabilities	593
4.19.6 Supported Resources for Each Type	597
4.19.7 Troubleshooting	599
4.19.8 Commands	600
4.20 Virtual Machine	607
4.20.1 Virtual Machine Configuration	607
4.20.2 Commands	610
4.21 Back-Up Battery Units	633
4.21.1 BBU Calibration Procedure	633
4.21.2 BBU Self-Test	634
4.21.3 BBU Shut-Off Timer	634
4.21.4 Commands	636
4.22 IP Table Filtering	644
4.22.1 Configuring IP Table Filtering	644
4.22.2 Modifying IP Table Filtering	646
4.22.3 Rate-limit Rule Configuration	646
4.22.4 Commands	647
4.23 Resource Scale	653
4.23.1 Ethernet Resources	653
4.23.2 Proxy-ARP Resources	655
4.23.3 Commands	657
Chapter 5 Ethernet Switching	661
5.1 Interface	661
5.1.1 Break-Out Cables	661
5.1.2 56GbE Link Speed	663
5.1.3 Transceiver Information	665

5.1.4	High Power Transceivers	665
5.1.5	Forward Error Correction	665
5.1.6	Commands	666
5.2	Interface Isolation	692
5.2.1	Configuring Isolated Interfaces	692
5.2.2	Commands	694
5.3	Link Aggregation Group (LAG)	700
5.3.1	Configuring Static Link Aggregation Group (LAG)	700
5.3.2	Configuring Link Aggregation Control Protocol (LACP)	700
5.3.3	Commands	702
5.4	MLAG	722
5.4.1	MLAG Keepalive and Failover	724
5.4.2	Unicast and Multicast Sync	724
5.4.3	MLAG Port Sync	724
5.4.4	MLAG Virtual System-MAC	725
5.4.5	Upgrading MLAG Pair	725
5.4.6	Configuring MLAG	725
5.4.7	Commands	730
5.5	VLANs	748
5.5.1	Configuring Access Mode and Assigning Port VLAN ID (PVID)	748
5.5.2	Configuring Hybrid Mode and Assigning Port VLAN ID (PVID)	749
5.5.3	Configuring Trunk Mode VLAN Membership	749
5.5.4	Configuring Hybrid Mode VLAN Membership	750
5.5.5	Commands	751
5.6	Voice VLAN	762
5.6.1	Configuring Voice VLAN	763
5.6.2	Limitations	766
5.7	QinQ	767
5.7.1	QinQ Operation Modes	767
5.7.2	Configuring QinQ	767
5.7.3	Commands	770
5.8	MAC Address Table	771
5.8.1	Configuring Unicast Static MAC Address	771
5.8.2	MAC Learning Considerations	771
5.8.3	Commands	772
5.9	Spanning Tree	779
5.9.1	Port Priority and Cost	779
5.9.2	Port Type	779
5.9.3	BPDU Filter	780
5.9.4	BPDU Guard	780
5.9.5	Loop Guard	780

5.9.6	Root Guard	781
5.9.7	MSTP	781
5.9.8	RPVST	781
5.9.9	Commands	784
5.10	OpenFlow	814
5.10.1	Flow Table	814
5.10.2	Configuring OpenFlow	815
5.10.3	Commands	817
5.11	OVS VTEP	828
5.11.1	Configuring OVS VTEP	828
5.11.2	Commands	830
5.12	IGMP Snooping	847
5.12.1	Configuring IGMP Snooping	847
5.12.2	Defining a Multicast Router Port on a VLAN	847
5.12.3	IGMP Snooping Querier	849
5.12.4	Commands	850
5.13	Link Layer Discovery Protocol (LLDP)	870
5.13.1	Configuring LLDP	870
5.13.2	DCBX	871
5.13.3	Commands	872
5.14	Quality of Service (QoS)	888
5.14.1	QoS Classification	888
5.14.2	QoS Rewrite	890
5.14.3	Queuing and Scheduling (ETS) for SwitchX	891
5.14.4	Queuing and Scheduling (ETS) for Spectrum	892
5.14.5	RED and ECN	894
5.14.6	Commands	896
5.15	Access Control List	928
5.15.1	Configuring Access Control List	928
5.15.2	ACL Actions	928
5.15.3	Commands	930
5.16	Port Mirroring	944
5.16.1	Mirroring Sessions	944
5.16.2	Configuring Mirroring Sessions	947
5.16.3	Verifying Mirroring Sessions	949
5.16.4	Commands	950
5.17	sFlow	960
5.17.1	Flow Samples	960
5.17.2	Statistical Samples	961
5.17.3	sFlow Datagrams	961
5.17.4	Sampled Interfaces	961
5.17.5	Configuring sFlow	961

5.17.6 Verifying sFlow	962
5.17.7 Commands	964
5.18 Transport Applications	976
5.18.1 RDMA over Converged Ethernet (RoCE)	976
5.19 802.1x Protocol	980
5.19.1 802.1x Operating Modes	980
5.19.2 Configuring 802.1x	981
5.19.3 Commands	982
5.20 Priority Flow Control	998
5.20.1 Flow Control Threshold Configuration for Spectrum	999
5.20.2 Commands	1001
5.21 Shared Buffers	1005
5.21.1 Packet Buffering Classification	1005
5.21.2 Buffering Allocation	1006
5.21.3 Pools	1007
5.21.4 Default Configurations	1007
5.21.5 Configuration Example	1008
5.21.6 Commands	1011
Chapter 6 IP Routing	1025
6.1 General	1025
6.1.1 IP Interfaces	1025
6.1.2 Equal Cost Multi-Path Routing (ECMP)	1028
6.1.3 Virtual Routing and Forwarding	1030
6.1.4 Commands	1031
6.2 OSPF	1072
6.2.1 Router ID	1072
6.2.2 ECMP	1072
6.2.3 Configuring OSPF	1073
6.2.4 Verifying OSPF	1074
6.2.5 Commands	1077
6.3 BGP	1115
6.3.1 State Machine	1115
6.3.2 Configuring BGP	1115
6.3.3 Verifying BGP	1117
6.3.4 Commands	1118
6.3.5 IP AS-Path Access-List	1171
6.3.6 IP Community-List	1173
6.4 Policy Rules	1176
6.4.1 Route Map	1176
6.4.2 IP Prefix-List	1207

6.5	Multicast (IGMP and PIM)	1210
6.5.1	Bidirectional PIM	1210
6.5.2	PIM Load-Sharing	1211
6.5.3	Bootstrap Router	1211
6.5.4	Configuring Multicast	1212
6.5.5	Commands	1215
6.6	VRRP	1258
6.6.1	Load Balancing	1258
6.6.2	Configuring VRRP	1259
6.6.3	Verifying VRRP	1261
6.6.4	Commands	1262
6.7	MAGP	1273
6.7.1	Configuring MAGP	1273
6.7.2	Commands	1275
6.8	DHCP Relay	1281
6.8.1	DHCP-R VRF Auto-Helper	1281
6.8.2	Commands	1282
Chapter 7	InfiniBand Switching	1294
7.1	Node Name	1294
7.1.1	Commands	1294
7.2	Fabric	1296
7.2.1	Commands	1296
7.3	IB Router	1301
7.3.1	Configuring IB Router	1302
7.3.2	Subnet Prefix Checking	1304
7.3.3	Commands	1305
7.4	Interface	1311
7.4.1	Transceiver Information	1311
7.4.2	High Power Transceivers	1311
7.4.3	Forward Error Correction	1311
7.4.4	Commands	1313
7.5	Subnet Manager (SM)	1336
7.5.1	Enabling Subnet Manager	1336
7.5.2	Partitions	1336
7.5.3	Adaptive Routing	1337
7.5.4	Commands	1338
7.6	Subnet Manager (SM) High Availability (HA)	1504
7.6.1	Joining, Creating or Leaving an InfiniBand Subnet ID	1505
7.6.2	MLNX-OS Management Centralized Location	1505
7.6.3	High Availability Node Roles	1505
7.6.4	Configuring MLNX-OS SM HA Centralized Location	1506

7.6.5	Creating and Adding Systems to an InfiniBand Subnet ID	1506
7.6.6	Restoring Subnet Manager Configuration	1507
7.6.7	Commands	1510
7.7	Fabric Inspector	1517
7.7.1	Running Diagnostics	1517
7.7.2	Mapping GUIDs to Node Names	1522
7.7.3	Importing ibdiagnet Fabric Data	1523
7.7.4	Commands	1524
Chapter 8	Gateway	1544
8.1	Proxy-ARP Prerequisites	1544
8.2	Proxy-ARP Overview	1545
8.2.1	Proxy-ARP Modes	1545
8.2.2	Proxy-ARP DHCP	1548
8.2.3	Proxy-ARP High Availability	1549
8.2.4	Proxy-ARP Interface	1553
8.2.5	Proxy-ARP HA Resources	1555
8.3	Proxy-ARP Event Notifications	1556
8.4	Proxy-ARP Configuration	1558
8.4.1	Proxy-ARP Mode Configuration	1558
8.4.2	Standalone Proxy-ARP Configuration	1558
8.4.3	High Availability Proxy-ARP Configuration	1561
8.5	Advanced Settings	1564
8.5.1	Default Gateway	1564
8.5.2	vTCA Interface	1564
8.5.3	MTU	1566
8.6	Commands	1567
8.6.1	Config	1567
8.6.2	Interface Proxy-ARP	1574
8.6.3	Show	1600
Appendix A	MEX6200 System	1622
A.1	MEX6200 Overview	1622
A.2	Getting Started	1622
A.3	Fault Management	1624
A.4	Alarms	1624
A.5	Events	1624
A.6	Alarm and Event Messages	1625
A.7	Configuration Management	1629
A.8	Performance Monitoring	1630
A.9	Native Signal	1630
A.10	Optical Level	1633

A.11	Maintenance	1636
A.12	Upgrading Software on the MEX6200	1637
Appendix B	Enhancing System Security According to NIST SP 800-131A .	1638
B.1	Overview	1638
B.2	Web Certificate	1638
B.3	SNMP	1639
B.4	SSH	1639
B.5	HTTPS	1640
B.6	LDAP	1641
Appendix C	Show Commands Supported by JSON API	1646

List of Tables

Table 1:	Reference Documents	43
Table 2:	Glossary	43
Table 3:	General System Features	47
Table 4:	Ethernet Features	48
Table 5:	InfiniBand Features	49
Table 6:	Serial Terminal Program Configuration for PPC Based Systems	52
Table 7:	Serial Terminal Program Configuration for x86 Based Systems	53
Table 8:	Configuration Wizard Session - IP Configuration by DHCP	54
Table 9:	Configuration Wizard Session - IP Zeroconf Configuration	56
Table 10:	Configuration Wizard Session - Static IP Configuration	57
Table 11:	MLNX-OS Licenses	63
Table 12:	CLI Modes and Config Context	73
Table 13:	Syntax Conventions	74
Table 14:	Angled Brackets Parameter Description	77
Table 15:	CLI Keyboard Shortcuts	81
Table 16:	WebUI Setup Submenus	83
Table 17:	WebUI System Submenus	84
Table 18:	WebUI Security Submenus	85
Table 19:	WebUI Ports Submenus	85
Table 20:	WebUI Status Submenus	86
Table 21:	WebUI IB SM Mgmt Submenus	87
Table 22:	WebUI Fabric Inspctr Submenus	88
Table 23:	WebUI ETH Mgmt Submenus	88
Table 24:	WebUI IP Route Submenus	89
Table 25:	Module Type	91
Table 26:	Device Type	91
Table 27:	Sensor Type	92
Table 28:	Supported Event Notifications and MIB Mapping	351
Table 29:	User Roles (Accounts) and Default Passwords	393
Table 30:	Chassis Manager Information	494
Table 31:	System Health Monitor Alerts Scenarios	495
Table 32:	LWR Configuration Behavior	502

Table 33: Standard MIBs – Textual Conventions and Conformance MIBs.	547
Table 34: Standard MIBs – Chassis and Switch	547
Table 35: Private MIBs Supported.	549
Table 36: SNMP Traps	549
Table 37: Supported SET OIDs.	554
Table 38: Ethernet and Port-Channel Interface Capabilities	593
Table 39: VLAN Capabilities.	593
Table 40: L2 Ethernet and Port-Channel Interface Capabilities	593
Table 41: LAG Capabilities	594
Table 42: L3 Interface Capabilities	594
Table 43: OSPF Interface Capabilities.	594
Table 44: OSPF Area Capabilities.	596
Table 45: Router OSPF Capabilities.	596
Table 46: Protocol Enable/Disable Capabilities	596
Table 47: Fetched Image Capabilities	596
Table 48: Installed Image Capabilities	597
Table 49: Fetched Image Capabilities	597
Table 50: Number of Resources per Node in Strict Mode	653
Table 51: Number of Resources per Node in Strict Mode for SwitchX Based Systems ...	654
Table 52: Number of Resources per Node in Strict Mode for Spectrum Based Systems ..	654
Table 53: Number of Resources per Node in Strict Proxy-ARP Unicast Mode	655
Table 54: Number of Resources per Node in Strict Proxy-ARP Multicast Mode.	655
Table 55: Maximum Number of Resources per Node in Proxy-ARP Loose Unicast Mode.	655
Table 56: Maximum Number of Resources per Node in Proxy-ARP Loose Multicast Mode	656
Table 57: Supported VLANs by RPVST per Switch System	782
Table 58: Packet Classification Rules	888
Table 59: Default QoS Configuration	889
Table 60: Default Shaper Configuration	893
Table 61: Mirroring Parameters	945
Table 62: List of Statistical Counters.	961
Table 63: IPoIB Forwarding	1546
Table 64: IPoIB Multicast Forwarding.	1547
Table 65: Proxy-ARP DHCP Linux Mode Application	1548
Table 66: Proxy-ARP DHCP Windows Mode Application.	1549
Table 67: Proxy-ARP Interface Attributes.	1554

Table 68: Supported Proxy-ARP Event Notifications	1556
Table 69: Alarm Messages	1625
Table 70: Configuration Event Messages	1627
Table 71: Other Event Messages	1628
Table 72: Configuration Options of the MEX6200	1629
Table 73: Link Port Performance Monitoring Tab Parameters	1632
Table 74: Link Port Performance Monitoring Tab Parameters	1635
Table 75: System Maintenance Options of the MEX6200	1636
Table 76: JSON API Show Commands	1646

List of Figures

Figure 1:	Managing an Ethernet Fabric Using MLNX-OS	49
Figure 2:	Managing an InfiniBand Software Using MLNX-OS	50
Figure 3:	Managing a	50
Figure 4:	Console Ports for CS75x0 Managed Systems	51
Figure 5:	Console Ports for SB7700 Managed Systems	51
Figure 6:	Console Ports for SX60xx/SX65xx Managed Systems	52
Figure 7:	MLNX-OS Login Window	60
Figure 8:	EULA Prompt	61
Figure 9:	Welcome Popup	62
Figure 10:	Display After Login	62
Figure 11:	No Licenses Installed	64
Figure 12:	Enter License Key(s) in Text Box	65
Figure 13:	Installed License	66
Figure 14:	WebUI	83
Figure 15:	Index Scheme	90
Figure 16:	SX65xx with Dual Management Modules	508
Figure 17:	SX60xx's LEDs	508
Figure 18:	JSON API WebUI Example	563
Figure 19:	Accepting an Agent Request through the Console	590
Figure 20:	Break-Out Cable	661
Figure 21:	Interface Isolation Example	692
Figure 22:	Basic MLAG Setup	722
Figure 23:	Basic MLAG Topology	725
Figure 24:	Tagging Voice Packets with a Different VLAN ID	762
Figure 25:	MAC Learning Disable Example Case	771
Figure 26:	RPVST Network Config	782
Figure 27:	RPVST and RSTP Cluster	783
Figure 28:	RED/ECN Drop Profiles	894
Figure 29:	Overview of Mirroring Functionality	944
Figure 30:	Mirror to Analyzer Mapping	944
Figure 31:	Header Format Options	947
Figure 32:	Mirroring Session	948
Figure 33:	sFlow Functionality Overview	960
Figure 34:	RoCEv2 and RoCE Frame Format Differences	976

Figure 35: RoCEv2 Protocol Stack	977
Figure 36: Xon/Xoff Configuration	1000
Figure 37: ECMP	1028
Figure 38: Multiple Hash Functions	1029
Figure 39: OSPF Basic Topology	1073
Figure 40: Basic BGP Configuration	1115
Figure 41: Common VRRP Configuration with Load Balancing	1259
Figure 42: Site-Local Unicast GID Format	1301
Figure 43: Host-to-Host IB Router Unicast Flow	1302
Figure 44: SM HA Subnet	1504
Figure 45: Gateway	1544
Figure 46: Basic Gateway Setup	1545
Figure 47: Unicast ARP Flow	1546
Figure 48: Multicast ARP Flow	1547
Figure 49: High Availability Proxy-ARP Interface	1550
Figure 50: Proxy-ARP Interface	1554
Figure 51: MetroX Connectivity to Switch	1622
Figure 52: MetroX Ports Tab Sidebar	1623
Figure 53: MEX6200 Item Buttons	1623
Figure 54: Sidebar Buttons	1624
Figure 55: Link Port Performance Monitoring Window	1630
Figure 56: Native Signal Performance Monitoring	1631
Figure 57: Optical Level Performance Monitoring	1634
Figure 58: MEX6200 Software Upgrade Webpage	1637

Document Revision History

Rev 4.60 – January 31, 2016

Added:

-
- the command “protocol nve” on page 830
- the command “interface nve” on page 831
- the command “shutdown” on page 832
- the command “vxlan source interface loopback” on page 833
- the command “nve mode only” on page 834
- the command “clear nve counters” on page 835
- the command “show interfaces nve” on page 836
- the command “show interfaces nve counters” on page 837
- the command “show interfaces nve flood” on page 838
- the command “show interfaces nve mac-address-table” on page 839
- the command “show interfaces nve mac-address-table local learned unicast” on page 840
- the command “show interfaces nve mac-address-table remote configured multicast” on page 841
- the command “show interfaces nve mac-address-table remote configured unicast” on page 842
- the command “show interfaces nve peers” on page 843
- the command “ovs ovssdb server” on page 844
- the command “ovs ovssdb manager remote” on page 845
- the command “ovs ovssdb server listen” on page 846
- Section 3.1.7.3, ““json-print” CLI Option,” on page 80
- Section 4.18.2, “JSON API,” on page 558
- Section 4.18.4.3, “JSON API Commands,” on page 587
- Section 4.11, “Telemetry,” on page 374 (Including subsections)
- the command “neighbor no-password” on page 1144
- the command “neighbor no-route-map” on page 1150
- the command “neighbor no-update-source” on page 1157
- the command “show snmp set-permission” on page 583
- the command “snmp-server enable set-permission” on page 571

- the command “show interfaces ethernet counters tc” on page 904
- the command “show interfaces ethernet counters pg” on page 905
- the command “show interfaces ethernet counters pfc prio” on page 906
- the command “show ip igmp snooping membership” on page 864
- the command “show lldp remote” on page 883
- the command “ip dhcp relay” on page 1282
- the command “ip dhcp relay instance” on page 1289
- the command “vrf” on page 1286
- the command “port” on page 1287
- the command “vrf-auto-helper” on page 1288
- Section 5.11, “OVS VTEP,” on page 828
- Appendix C, “Show Commands Supported by JSON API,” on page 1646

Updated:

- Section 8.2.5, “Proxy-ARP HA Resources,” on page 1555
- Section 8.4.2, “Standalone Proxy-ARP Configuration,” on page 1558
- the command “show proxy-arp mode” on page 1614
- the command “system resource table” on page 657
- the command “show system resource table” on page 658
- the command “bestpath as-path multipath-relax” on page 1123
- the command “bgp listen range peer-group” on page 1126
- the command “cluster-id” on page 1127
- the command “graceful-restart helper” on page 1131
- the command “neighbor advertisement-interval” on page 1133
- the command “neighbor allowas-in” on page 1134
- the command “neighbor ebgp-multihop” on page 1136
- the command “neighbor export-localpref” on page 1137
- the command “neighbor import-localpref” on page 1138
- the command “neighbor local-as” on page 1139
- the command “neighbor maximum-prefix” on page 1140
- the command “neighbor next-hop-peer” on page 1141
- the command “neighbor peer-group” on page 1145
- the command “neighbor route-reflector-client” on page 1151
- the command “neighbor send-community” on page 1152

- the command “neighbor timers” on page 1154
- the command “neighbor transport connection-mode passive” on page 1155
- the command “router-id” on page 1161
- the command “match as-path” on page 1183
- the command “set community-list” on page 1197
- the command “hostname” on page 180
- the command “image options” on page 267 “serve all” parameter description
- Note in Section 4.6.5, “Downgrading MLNX-OS Software,” on page 248
- Section 4.7.4.2, “Text Configuration Files,” on page 272
- the command “show running-config” on page 304
- Section 4.18.3, “XML API,” on page 563
- Section 4.23, “Resource Scale,” on page 653
- the command “fec-override” on page 676
- the command “ip igmp snooping static-group” on page 856
- the command “ip igmp snooping version” on page 858
- the command “openflow mode hybrid” on page 819
- the command “show ip igmp snooping groups” on page 862
- the command “show ip igmp snooping querier” on page 866
- the command “show ip igmp snooping statistics” on page 868
- the command “show sflow” on page 975
- Section 6.3.4, “Commands,” on page 1118 under BGP
- Section 6.4.1, “Route Map,” on page 1176
- the command “address” on page 1283
- the command “always-on” on page 1284
- the command “information option” on page 1285
- the command “clear ip dhcp relay counters” on page 1290
- the command “show ip dhcp relay” on page 1292
- the command “show ip dhcp relay counters” on page 1292
- the command “clear ip bgp” on page 1119
- the command “ib sm log-max-size” on page 1359
- the command “show proxy-arp mode” on page 1614

Rev 4.50 – September 30, 2016

Added:

- Section 3.1.7.2, ““watch” CLI Monitoring Option,” on page 79
- the command “ntp server trusted-enable” on page 218
- the command “logging <syslog IP address> port” on page 307
- the command “install-from-usb” on page 616
- the command “boot-delay” on page 667
- the command “show interfaces ethernet rates” on page 684
- the command “show interfaces ethernet transceiver diagnostics” on page 689
- the command “show mac-address-table summary” on page 778
- Mellanox Community post to Section 7.3, “IB Router,” on page 1301
- the command “show interfaces ib transceiver diagnostics” on page 1332
- the command “ib sm m-key” on page 1365
- the command “ib sm mkey-lease” on page 1366
- the command “ib sm mkey-lookup” on page 1367
- the command “ib sm mkey-protect-level” on page 1368
- the command “ib sm rtr-aguid-enable” on page 1381
- the command “ib sm rtr-pr-flow-label” on page 1382
- the command “ib sm rtr-pr-mtu” on page 1383
- the command “ib sm rtr-pr-sl” on page 1385
- the command “ib sm rtr-pr-tclass” on page 1386
- the command “ib sm subnet-prefix-override” on page 1395
- the command “ib sm virt-default-hop-limit” on page 1404
- the command “ib sm virt-max-ports-in-process” on page 1405
- the command “show ib sm m-key” on page 1438
- the command “show ib sm mkey-lookup” on page 1439
- the command “show ib sm mkey-protect-level” on page 1440
- the command “show ib sm rtr-aguid-enable” on page 1453
- the command “show ib sm rtr-pr-flow-label” on page 1454
- the command “show ib sm rtr-pr-mtu” on page 1455
- the command “show ib sm rtr-pr-rate” on page 1456
- the command “show ib sm rtr-pr-sl” on page 1457
- the command “show ib sm subnet-prefix-override” on page 1466

- the command “show ib sm virt-default-hop-limit” on page 1474
- the command “show ib sm virt-max-ports-in-process” on page 1475
- Section 8.3, “Proxy-ARP Event Notifications,” on page 1556
- Appendix C, “Mellanox NEO™ on Switch” on page 1601

Updated:

- Section 1.2, “Ethernet Features,” on page 48 with Spectrum™ unicast addresses
- Section 4.7.4.1, “BIN Configuration Files,” on page 271
- Section 4.7.4.2, “Text Configuration Files,” on page 272
- the command “reset factory” on page 279
- the command “show running-config” on page 304
- Section 4.8.2, “Remote Logging,” on page 306 with Step 3
- Section 4.10.1, “Supported Events,” on page 351
- the command “username” on page 397
- Section 4.17.5, “Unit Identification LED,” on page 507
- the command “led uid” on page 515
- the command “show leds” on page 526
- Section 4.18.1.2, “Private MIB,” on page 549
- Section “Changing Configuration with SNMP” on page 556
- Section 4.20.1, “Virtual Machine Configuration,” on page 607
- the command “show interfaces ethernet” on page 678
- Section 5.4.6, “Configuring MLAG,” on page 725
- Section 5.9.7, “MSTP,” on page 781
- the command “controller-ip” on page 820
- Section 5.12, “IGMP Snooping,” on page 847 with IGMPv3 note
- Section 5.15.2, “ACL Actions,” on page 928
- the command “deny/permit (IPv4 TCP/UDP/ICMP ACL rule)” on page 935
- the command “monitor session” on page 950
- Section 6.7.1, “Configuring MAGP,” on page 1273
- Section 7.5.1, “Enabling Subnet Manager,” on page 1336
- the command “ib sm subnet-prefix” on page 1394
- Section A.4, “Configuration Management,” on page 1629

Rev 4.40 – June 28, 2016

Added:

- Section 4.6.5.3, “Switching to Partition with Older Software Version,” on page 251 for clarity
- Section 4.7.4.1, “BIN Configuration Files,” on page 271
- Section 4.14.1, “System File Encryption,” on page 445
- Section 4.17.5, “Unit Identification LED,” on page 507
- the command “crypto encrypt-data” on page 446
- the command “show crypto encrypt-data” on page 461
- the command “led uid” on page 515
- the command “show leds” on page 526
- the command “show protocols” on page 531
- the command “show system capabilities” on page 533
- Section 4.21.3, “BBU Shut-Off Timer,” on page 634
- the command “battery-backup-unit shut-off-timer” on page 640
- Section 5.1.5, “Forward Error Correction,” on page 665
- the command “show interfaces ethernet transceiver counters” on page 687
- the command “show interfaces ethernet transceiver counters details” on page 688
- the command “show interfaces ethernet transceiver raw” on page 691
- Section 5.2, “Interface Isolation,” on page 692
- the command “show interfaces port-channel counters” on page 718
- “Enabling L3 Forwarding with User VRF” on page 729
- the command “show interfaces mlag-port-channel counters” on page 745
- the command “switchport voice” on page 760
- Section 5.6, “Voice VLAN,” on page 762
- the command “ip igmp snooping clear counters” on page 853
- the command “ip igmp snooping version” on page 858
- the command “show ip igmp snooping querier counters” on page 867
- the command “lldp med-tlv-select” on page 879
- Section 5.14.1, “QoS Classification,” on page 888
- Section 5.14.2, “QoS Rewrite,” on page 890
- Section 5.14.3, “Queuing and Scheduling (ETS) for SwitchX,” on page 891
- Section 5.14.6.1, “QoS Classification,” on page 896

- the command “no area” on page 1089
- the command “neighbor send-community” on page 1152
- Section 7.3, “IB Router,” on page 1301
- Section 7.4.3, “Forward Error Correction,” on page 1311
- the command “switchport access subnet” on page 1323
- the command “show interfaces ib transceiver raw” on page 1334

Updated:

- the command “Starting the Web User Interface (WebUI)” on page 59
- the command “license delete” on page 70
- the command “ip arp timeout” on page 192
- the command “system profile” on page 240
- Section 4.6, “Software Management,” on page 242 with note about Switch-IB™ 2 interoperability
- Section 4.6.7, “Image Maintenance via Mellanox ONIE,” on page 254
- the command “image fetch” on page 262
- the command “configuration switch-to” on page 294
- “Changing Configuration with SNMP” on page 556 with BinaryDelete and TextDelete commands
- the command “show asic-version” on page 520
- the command “power enable” on page 516
- the command “show inventory” on page 525
- the command “show interfaces ib internal notification” on page 584
- the command “show battery-backup-unit” on page 642
- the command “show interfaces ethernet” on page 678
- the command “show interfaces ethernet counters” on page 681
- the command “show interfaces port-channel” on page 716
- the command “show interfaces mlag-port-channel” on page 744
- Section 5.8.1, “Configuring Unicast Static MAC Address,” on page 771
- the command “show openflow detail” on page 823
- the command “show openflow flows” on page 824
- the command “show openflow statistics” on page 825
- the command “ip igmp snooping static-group” on page 856
- the command “show ip igmp snooping” on page 861

- the command “show ip igmp snooping groups” on page 862
- the command “show ip igmp snooping statistics” on page 868
- Section 5.13.2, “DCBX,” on page 871
- the command “show lldp interfaces” on page 882
- Section 5.21, “Shared Buffers,” on page 1005
- the command “ip arp timeout” on page 1063
- the command “router ospf” on page 1078
- the command “area stub” on page 1087
- the command “show ip ospf” on page 1104
- the command “show ip ospf border-routers” on page 1105
- the command “show ip ospf database” on page 1106
- the command “show ip ospf interface” on page 1108
- the command “show ip ospf neighbors” on page 1110
- the command “neighbor peer-group” on page 1145
- the command “always-on” on page 1284
- the command “clear ip dhcp relay counters” on page 1290
- the command “ip dhcp relay information option circuit-id” on page 1291
- the command “show ip dhcp relay” on page 1292
- the command “show ip dhcp relay counters” on page 1292
- the command “show guides” on page 1298
- the command “show lids” on page 1300
- the command “show interfaces ib” on page 1324
- the command “show interfaces ib status” on page 1326

Removed “Security Vulnerabilities and Exposures” appendix and moved it to www.mellanox.com/page/mlnx_os_security_vulnerabilities_and_exposures

Rev 4.30 – March 02, 2016

Added:

- Section 3.1.8, “CLI Shortcuts,” on page 81
- Section 3.4, “Management Information Bases (MIBs),” on page 90
- the command “cli max-sessions” on page 97
- the command “show cli max-sessions” on page 104
- the command “show cli num-sessions” on page 105

- the command “banner logout” on page 109
- the command “banner logout-local” on page 110
- the command “banner logout-remote” on page 111
- the command “ssh server login attempts” on page 118
- the command “ssh server login timeout” on page 119
- Section 4.1.6, “Configuring Hostname via DHCP (DHCP Client Option 12),” on page 154
- the command “dhcp hostname” on page 167
- Section 4.2.1, “NTP Authenticate,” on page 206
- Section 4.2.2, “NTP Authentication Key,” on page 206
- the command “ntp authenticate” on page 211
- the command “ntp authentication-key” on page 212
- the command “ntp peer disable” on page 213
- the command “ntp peer keyID” on page 214
- the command “ntp peer version” on page 215
- the command “ntp server disable” on page 216
- the command “ntp server keyID” on page 217
- the command “ntp server version” on page 219
- the command “ntp trusted-key” on page 220
- the command “show ntp configured” on page 223
- the command “show ntp keys” on page 224
- Section 4.6.7, “Image Maintenance via Mellanox ONIE,” on page 254
- Section 4.13.2.1, “User Re-authentication,” on page 394
- the command “aaa authentication attempts fail-delay” on page 404
- the command “show system type” on page 536
- Section 4.22, “IP Table Filtering,” on page 644
- the command “show interfaces ethernet counters” on page 681
- Section 5.14.5, “RED and ECN,” on page 894
- Section 5.14.6.4, “RED & ECN,” on page 925
- the command “ip multicast filter” on page 1599
- the command “show ip multicast interface proxy-arp” on page 1615
- the command “show ip multicast interface proxy-arp count” on page 1616
- the command “show ip multicast filter interface proxy-arp” on page 1618

Updated:

- the command “show banner” on page 113
- the command “ssh server login attempts” on page 118
- the command “ssh server security strict” on page 122
- the command “show ssh server” on page 131
- Section 4.1, “Management Interface,” on page 152
- the command “show interface configured” on page 179
- the command “show ntp” on page 222
- the command “show aaa authentication attempts” on page 416
- Table 31, “System Health Monitor Alerts Scenarios,” on page 495
- the command “show inventory” on page 525
- Section 4.21.1, “BBU Calibration Procedure,” on page 633 with section
- the command “battery-backup-unit calibrate-battery” on page 638 note section
- the command “battery-backup-unit calibrate-battery foreground” on page 639 with note
- the command “show battery-backup-unit details” on page 643
- the command “show power” on page 529
- the command “show power consumers” on page 530
- Table 34, “Standard MIBs – Chassis and Switch,” on page 547
- Section 5.1.1, “Break-Out Cables,” on page 661
- Section 5.1.2, “56GbE Link Speed,” on page 663
- Section 5.1.5, “Forward Error Correction,” on page 665
- the command “speed” on page 673
- the command “deny/permit (MAC ACL rule)” on page 932
- the command “deny/permit (IPv4 ACL rule)” on page 934
- the command “deny/permit (IPv4 TCP/UDP/ICMP ACL rule)” on page 935
- Section 5.16.1.2, “Destination Interface,” on page 946
- the command “add source interface” on page 953
- the command “header-format” on page 954
- the command “show monitor session” on page 958
- Section 5.21, “Shared Buffers,” on page 1005
- the command “ip load-sharing” on page 1058
- the command “ip arp timeout” on page 1063
- the command “bgp listen range peer-group” on page 1126

- the command “show ip mroute” on page 1243
- the command “vrrp” on page 1263
- the command “pkey” on page 1478
- the command “defmember” on page 1479
- the command “member” on page 1480
- the command “ipoib” on page 1481
- the command “mtu” on page 1482
- the command “ib sm log-max-size” on page 1359
- the command “ib sm virt” on page 1403
- the command “rate” on page 1483
- the command “scope” on page 1484
- the command “sl” on page 1485
- Section 8.2.2, “Proxy-ARP DHCP,” on page 1548
- Section A.6.1, “Upgrading Software on the MEX6200,” on page 1637
- Section B.4, “SSH,” on page 1639

Rev 4.20 – August 16, 2015

Added:

- Section 4.6.7, “Image Maintenance via Mellanox ONIE,” on page 254
- Section 4.13.3, “System Secure Mode,” on page 395
- the command “system secure-mode enable” on page 443
- the command “show system secure-mode” on page 444
- Section 4.17.2.1, “Power Supply Options,” on page 501
- the command “show asic-version” on page 520
- the command “switchport dot1q-tunnel qos-mode” on page 756
- Section 5.7, “QinQ,” on page 767
- the command “dot1x host-mode” on page 985
- the command “show ip route” on page 1059
- the command “vlan-pop” on page 939
- the command “vlan-push” on page 940
- Section 8.2.2, “Proxy-ARP DHCP,” on page 1548

Updated:

- Table 8, “Configuration Wizard Session - IP Configuration by DHCP,” on page 54

- Section 2.4, “Licenses,” on page 63
- the command “ssh server host-key” on page 115
- Table 28, “Supported Event Notifications and MIB Mapping,” on page 351
- notes of the command “aaa authorization” on page 413
- Table 31, “System Health Monitor Alerts Scenarios,” on page 495
- the command “show module” on page 528
- the command “snmp-server user” on page 579
- Section 5.1.2, “56GbE Link Speed,” on page 663
- the command “switchport mode” on page 754
- the command “ip ospf authentication-key” on page 1101
- the command “neighbor password” on page 1143
- the command “neighbor peer-group” on page 1145
- the command “interface ib internal notification” to the commands:
 - “interface ib internal notification link-speed-mismatch” on page 1321
 - “interfaces ib internal notification link-state-change” on page 1322
- Section 7.6.1, “Joining, Creating or Leaving an InfiniBand Subnet ID,” on page 1505
- Section 8.2.3, “Proxy-ARP High Availability,” on page 1549 with DHCP note

Rev 4.10 – June 11, 2015

Added:

- Section 2.1, “Configuring the Switch for the First Time,” on page 51 with MLNX-OS® Boot Menu step
- the command “ssh server security strict” on page 122
- the command “ssh server tcp-forwarding enable” on page 123
- Section 4.1.5, “In-Band Management,” on page 153
This feature can now be enabled with IP Routing. Also updated the flow of setting an in-band management channel.
- the command “interface ib internal phy-profile enable llr64” on page 230
- the command “show module” on page 528
- Section 4.21.1, “BBU Calibration Procedure,” on page 633
- Section 4.21.2, “BBU Self-Test,” on page 634
- the command “battery-backup-unit calibrate-battery” on page 638
- the command “battery-backup-unit calibrate-battery foreground” on page 639
- the command “battery-backup-unit test-battery” on page 641

- the command “show battery-backup-unit” on page 642
- the command “show battery-backup-unit details” on page 643
- Section 5.1.1, “Break-Out Cables,” on page 661
- the command “ip address dhcp” on page 675
- the command “ip address dhcp” on page 709
- Section 5.4.4, “MLAG Virtual System-MAC,” on page 725
- Section 5.4.5, “Upgrading MLAG Pair,” on page 725
- Section 5.19, “802.1x Protocol,” on page 980
- Section 6.1.3, “Virtual Routing and Forwarding,” on page 1030
- the command “ip l3” on page 1031
- the command “vrf definition” on page 1032
- the command “routing-context vrf” on page 1033
- the command “description” on page 1035
- the command “rd” on page 1036
- the command “vrf forwarding” on page 1037
- the command “show routing-context vrf” on page 1039
- the command “show vrf” on page 1040
- the command “ip address dhcp” on page 1045
- Section 6.2, “IPv6,” on page 990 commands by adding loopback interface configuration mode to the commands
- Section 6.5.2, “PIM Load-Sharing,” on page 1211
- the command “ip pim multipath rp” on page 1229
- the command “ib sm drop-event-subscription” on page 1343
- the command “ib sm virt” on page 1403
- the command “show interfaces ib internal” on page 1327
- the command “show interfaces ib internal capabilities” on page 1328
- the command “show interfaces ib internal llr” on page 1329
- the command “show interfaces ib internal status” on page 1330

Updated:

- the command “tcpdump” on page 205
- Section 4.6.1, “Upgrading MLNX-OS Software,” on page 242 with HA group note
- Section 4.6.3, “Upgrading MLNX-OS HA Groups,” on page 247
- the command “show inventory” on page 525

- the command “show asic-version” on page 520
- the notes in command “battery-backup-unit calibrate-battery” on page 638
- Section 5.4.1, “MLAG Keepalive and Failover,” on page 724
- Step 10 in Section 5.4.6, “Configuring MLAG,” on page 725
- the example of the command “upgrade-timeout” on page 741
- the command “ip routing” on page 1034
- the command “show ip routing” on page 1038
- the command “show ip interface” on page 1052
- the command “interface loopback” on page 1053 “id” parameter range
- the command “ip route” on page 1057
- the command “show ip route” on page 1059
- the command “clear ip arp” on page 1064
- the command “show ip arp” on page 1065
- the command “ping” on page 1067
- the command “traceroute” on page 1068
- the command “tcpdump” on page 1070
- the command “show guids” on page 1298
- the command “show lids” on page 1300
- the command “interface ib” on page 1313
- the command “speed” on page 1317

Removed:

- the command “interface vlan create” from Section 4.1.7, “Commands,” on page 155
- the command “ipv6 dhcp client”

Split:

- the command “ipv6 dhcp”
- the command “show {guids | system guid}”

Rev 3.70 – March 19, 2015

Updated:

- the command “speed” on page 1317
- the command “show interfaces ib” on page 1324
- the command “show interfaces ib status” on page 1326
- the command “ib sm force-link-speed” on page 1346

- the command “[show ib sm force-link-speed](#)” on page 1413
- the command “[show ib sm force-link-speed-ext](#)” on page 1414

Rev 3.70 – March 19, 2015

No changes

Rev 3.60 – March 05, 2015

Added:

- MLAG configuration [Step 10](#)
- the command “[system-mac](#)” on page 740
- the command “[upgrade-timeout](#)” on page 741
- [Section 5.9.4, “BPDU Guard,”](#) on page 780

Updated:

- MLAG configuration verification [Step 1](#) with system MAC and upgrade timeout
- the command “[show mlag](#)” on page 742

Rev 3.60 – March 05, 2015

No changes

Rev 3.50 – February 24, 2015

Added:

- the command “[show version concise](#)” on page 540

Updated:

- the command “[show uboot](#)” on page 538

Rev 3.40 – February 11, 2015

Added:

- “List of Tables” and “List of Figures” Sections
- Updated [Section 2.4, “Licenses,”](#) on page 63
- the command “[license delete](#)” on page 70
- the command “[license install](#)” on page 71
- the command “[telnet](#)” on page 132
- the command “[terminal](#)” on page 101
- the command “[web cache-enable](#)” on page 137

- the command “ip default-gateway” on page 158
- the command “boot system” on page 259
- the command “configuration write” on page 299
- the command “logging trap” on page 326
- the command “email autosupport enable” on page 355
- the command “crypto ipsec ike” on page 447
- Section 4.21, “Back-Up Battery Units,” on page 633
- the command “lACP-individual enable” on page 708
- the command “show interfaces port-channel” on page 716
- the command “show interfaces port-channel compatibility-parameters” on page 719
- the command “show interfaces port-channel load-balance” on page 720
- the command “show interfaces port-channel summary” on page 721
- Section 5.9.8, “RPVST,” on page 781
- the command “spanning-tree vlan forward-time” on page 803
- the command “spanning-tree vlan hello-time” on page 804
- the command “spanning-tree vlan max-age” on page 805
- the command “spanning-tree vlan priority” on page 806
- the command “show spanning-tree vlan” on page 812
- Section 6.2, “IPv6,” on page 990
- the command “auto-cost reference-bandwidth” on page 1081
- the command “ib sm allow-both-pkeys” on page 1340
- the command “show ip multicast interface proxy-arp” on page 1615

Updated:

- Section 2.3, “Starting the Web User Interface (WebUI),” on page 59
- the command “image options” on page 267
- the command “reload” on page 278
- Section 4.8.2, “Remote Logging,” on page 306
- the command “logging debug-files” on page 311
- Section 4.9.1, “Commands,” on page 331
- Section 4.13.1, “User Accounts,” on page 393
- the command “username” on page 397
- the command “aaa authentication attempts fail-delay” on page 404
- the command “radius-server host” on page 418

- the command “tacacs-server host” on page 422
- Table 31, “System Health Monitor Alerts Scenarios,” on page 495
- the command “snmp-server auto-refresh” on page 564
- the command “snmp-server user” on page 579
- the command “show interfaces ethernet description” on page 683
- the command “show interfaces ethernet status” on page 685
- the command “show interfaces port-channel summary” on page 721
- the command “show interfaces mlag-port-channel summary” on page 746
- the command “spanning-tree mode” on page 785
- the command “show spanning-tree” on page 807
- the command “show spanning-tree detail” on page 808
- the command “show spanning-tree interface” on page 809
- the command “show spanning-tree mst” on page 810
- the command “show spanning-tree root” on page 811
- Section 5.12.2, “Defining a Multicast Router Port on a VLAN,” on page 847
- the command “dcb application-priority” on page 880
- the command “dcb priority-flow-control enable” on page 1001
- Section 5.17.1, “Flow Samples,” on page 960
- the command “ip arp timeout” on page 1063
- the command “redistribute” on page 1083
- Section 7.4.2, “High Power Transceivers,” on page 1311
- the command “ib sm root-guid” on page 1379
- the command “defmember” on page 1479
- the command “member” on page 1480
- the command “rate” on page 1483
- the command “ib qos level” on page 1495
- Chapter 8, “Gateway” on page 1544
- the command “show interfaces proxy-arp ha multicast-list” on page 1617

Rev 3.30 – November 19, 2014

Added:

- Section 5.1.4, “High Power Transceivers,” on page 665
- Section 7.4.2, “High Power Transceivers,” on page 1311

Updated:

- the command “web https” on page 145
- the command “show interfaces ethernet” on page 678
- the command “show interfaces ethernet transceiver” on page 686
- the command “dcb application-priority” on page 880
- the command “show interfaces ib” on page 1324
- the command “show interfaces ib transceiver” on page 1331
- Section B.5, “HTTPS,” on page 1640
- Section B.7, “Password Hashing,” on page 1757

Rev 3.20 – November 09, 2014

Added:

- MAC addresses note in Section 4.4, “Virtual Protocol Interconnect (VPI),” on page 234
- Section 4.20, “Virtual Machine,” on page 607
- Section 5.8.2, “MAC Learning Considerations,” on page 771
- the command “mac-learning disable” on page 774
- Appendix B, “Enhancing System Security According to NIST SP 800-131A,” on page 1638

Updated:

- Section 1.2, “Ethernet Features,” on page 48
- Section 3.2, “Web Interface Overview,” on page 82
- the command “reset factory” on page 279
- Section 4.18.1.7, “SNMP SET Operations,” on page 553
- the command “interface port-channel” on page 702
- the command “show lacp interfaces neighbor” on page 712
- Section 5.4, “MLAG,” on page 722
- the command “mlag-channel-group mode” on page 737
- the command “show mlag statistics” on page 747
- the command “ip icmp redirect” on page 1051
- Section 6.3, “BGP,” on page 1115
- Section 6.6.2, “Configuring VRRP,” on page 1259
- the command “show fabric” on page 1297
- the command “proxy-arp ha” on page 1571

- Appendix A, “MEX6200 System,” on page 1622

Replaced:

- the command “show lacp interfaces port-channel” with the command “show lacp” on page 714
- the command “show lacp system-identifier” with the command “show lacp interfaces system-identifier” on page 715

Rev 3.10 – July 20, 2014

Added:

- Section 5.18, “Transport Applications,” on page 976
- Section 6.1.1, “IP Interfaces,” on page 1025
- Section 6.3, “BGP,” on page 1115
- the command “show ip pim upstream joins” on page 1239

Updated:

- Chapter 1, “Introduction” on page 47
- Section 4.18.1.8, “IF-MIB and Interface Information,” on page 558
- Section 4.18.3, “XML API,” on page 563
- MAC addresses note in Section 5.4, “MLAG,” on page 722
- Chapter 6, “IP Routing” on page 1025 with the appropriate configuration modes for the new configuration contexts and commands added
- the command “route-map” on page 1177
- the command “continue <sequence-number>” on page 1178
- the command “abort” on page 1180
- the command “exit” on page 1181
- Section 6.5, “Multicast (IGMP and PIM),” on page 1210
- the command “ip pim join-prune-interval” on page 1226
- the command “show ip pim bsr” on page 1232
- the command “show ip mroute” on page 1243
- CPU type note in Section 7.6, “Subnet Manager (SM) High Availability (HA),” on page 1504

Rev 3.00 – June 05, 2014

Updated:

- Section 6.5, “Multicast (IGMP and PIM),” on page 1210
- Section 6.6.3, “Verifying VRRP,” on page 1261

Rev 2.90 – 19 May, 2014

Added:

- Section 4.18.1.8, “IF-MIB and Interface Information,” on page 558
- Section 6.5, “Multicast (IGMP and PIM),” on page 1210

Updated:

- the command “port type” on page 236
- the command “show configuration” on page 302
- the command “show uboot” on page 538
- the command “show voltage” on page 541
- Section 5.4, “MLAG,” on page 722
- the command “show mlag” on page 742
- Section 6.1.4.2, “IP Interfaces,” on page 1041
- Section 6.1.4.4, “Loopback Interface,” on page 1053
- Section 8.2.3.4, “Proxy-ARP Load Balancing,” on page 1552

Rev 2.80 – May 08, 2014

Added:

- supported versions note in Section 5.12, “IGMP Snooping,” on page 847
- Section 6.6, “VRRP,” on page 1258
- Section 6.7, “MAGP,” on page 1273
- Section 6.8, “DHCP Relay,” on page 1281

Rev 2.70 – April 30, 2014

Added:

- Appendix B, “Enhancing System Security According to NIST SP 800-131A,” on page 1638
- supported versions note in Section 5.12, “IGMP Snooping,” on page 847

Updated:

- the command “show ssh server” on page 131
- the command “web auto-logout” on page 135
- the command “web https” on page 145
- the command “show web” on page 151
- the command “show usernames” on page 399
- the command “ldap base-dn” on page 425

- the command “[ldap ssl](#)” on page 436

Rev 2.60 – April 10, 2014

Updated:

- Table 35, “Private MIBs Supported,” on page 549

Rev 2.50 – April 2014

Updated:

- Section 3.1.7, “CLI Pipeline Operator Commands,” on page 78
- Section 4.17.7, “System Reboot,” on page 511
- the command “[show protocols](#)” on page 531
- the command “[show mac-address-table](#)” on page 776
- the command “[deny/permit \(MAC ACL rule\)](#)” on page 932
- the command “[show mac/ipv4 access-lists](#)” on page 942
- the command “[ha member ip address](#)” on page 1595

Added:

- Section 5.4, “MLAG,” on page 722
- configuration mode Config Interface MLAG Port Channel to the following commands:
 - “[flowcontrol](#)” on page 668
 - “[mtu](#)” on page 670
 - “[shutdown](#)” on page 671
 - “[description](#)” on page 672
 - “[speed](#)” on page 673
 - “[load-interval](#)” on page 674
 - “[clear counters](#)” on page 677
 - “[switchport mode](#)” on page 754
 - “[switchport access](#)” on page 757
 - “[spanning-tree port-priority](#)” on page 789
 - “[spanning-tree cost](#)” on page 790
 - “[spanning-tree port type](#)” on page 791
 - “[spanning-tree guard](#)” on page 792
 - “[ip igmp snooping fast-leave](#)” on page 854
 - “[dcb priority-flow-control mode on](#)” on page 1003
 - “[ipv4/mac port access-group](#)” on page 931

- “sflow enable (interface)” on page 974

Rev 2.40 – February, 2014

Updated:

- Section 4.6.6.3, “Importing Firmware and Changing the Default Firmware,” on page 253 – updated Step 1
- the command “show running-config” on page 304
- the command “show log” on page 328
- Section 4.14, “Cryptographic (X.509, IPSec) and Encryption,” on page 445
- Section 5.3.1, “Configuring Static Link Aggregation Group (LAG),” on page 700 – removed unnecessary step
- the command “lldp tlv-select” on page 878
- Chapter 8, “Gateway” on page 1544

Added:

- Section 3.1.7, “CLI Pipeline Operator Commands,” on page 78
- FCoE and SX1700 GW license in Section 2.4, “Licenses,” on page 63
- Section 4.18.1.8, “IF-MIB and Interface Information,” on page 558
- Section 4.17.7, “System Reboot,” on page 511

Rev 2.30 – January, 2014

Updated:

- Section 4.19.4, “Writing Configuration Classes,” on page 590
- the command “crypto certificate generation” on page 453
- the command “crypto certificate name” on page 455

Rev 2.20 – January, 2014

Updated:

- Section 4.19.5.11, “Installed Image Capabilities,” on page 597

Rev 2.10 – January, 2014

Added:

- Section 4.17.2.2, “Width Reduction Power Saving,” on page 502

Updated:

- Section 2.2, “Starting the Command Line (CLI),” on page 58

- Section 2.3, “Starting the Web User Interface (WebUI),” on page 59
- the command “system profile” on page 240
- Section 4.6.1, “Upgrading MLNX-OS Software,” on page 242 with EULA note
- Section 4.6.2, “Upgrading MLNX-OS Software on Director Switches,” on page 246 with a note
- Section 4.19, “Puppet Agent,” on page 589
- the command “load-interval” on page 674 with Config Interface Port Channel
- the command “spanning-tree port-priority” on page 789 with Config Interface Port Channel
- Section 5.10, “OpenFlow,” on page 814
- the command “openflow description (SwitchX)” on page 818
- the command “show openflow” on page 827
- the command “switchport {hybrid, trunk} allowed-vlan” on page 758 with Config Interface Port Channel
- the command “spanning-tree cost” on page 790 with Config Interface Port Channel
- the command “spanning-tree port type” on page 791 with Config Interface Port Channel
- the command “spanning-tree guard” on page 792 with Config Interface Port Channel
- the command “spanning-tree bpdupfilter” on page 793 with Config Interface Port Channel
- the command “deny/permit (IPv4 ACL rule)” on page 934
- the command “sflow enable (interface)” on page 974 with Config Interface Port Channel
- Section 6.2, “OSPF,” on page 1072
- the command “router-id” on page 1079
- Section 7.5.3, “Adaptive Routing,” on page 1337

Merged sections “Restoring Factory Default Configuration on a Switch System (Single Management Module)” and “Restoring Factory Default Configuration on a Switch Directors (Dual Management Modules)” under Section 4.7.3, “Restoring Factory Default Configuration,” on page 271.

Rev 2.00 – December 2013

Added:

- Section 5.1.3, “Transceiver Information,” on page 665
- Section 7.4.1, “Transceiver Information,” on page 1311
- the command “run-interval” on page 603
- a note to Section 8.5.3, “MTU,” on page 1566

Updated:

- Section 4.4, “Virtual Protocol Interconnect (VPI),” on page 234 with SX1012 and SX6012
- Section 4.6.1, “Upgrading MLNX-OS Software,” on page 242
- Section 4.6.4, “Deleting Unused Images,” on page 248
- Section 4.9, “Debugging,” on page 330
- the example of the command “show cpld” on page 522
- “Notification Indicator” column in Section 8.4.2, “Standalone Proxy-ARP Configuration,” on page 1558
- the command “show puppet-agent” on page 605
- the command “lldp tlv-select” on page 878
- the command “interface proxy-arp” on page 1570
- the command “ip address” on page 1574
- the command “ip vlan” on page 1576
- the command “ip pkey” on page 1577
- the command “ip route” on page 1578
- the command “counters” on page 1583

Moved:

Section 3.3, “Secure Shell (SSH),” on page 89 from 4.13.2

Removed:

- mention of the MLNX-OS Command Reference Guide
- the command “lldp tlv-select dcbx”

Rev 1.90 – November 2013

Added Appendix A, “MEX6200 System,” on page 1622

Rev 1.80 – October 2013

Added:

- Section 4.19, “Puppet Agent,” on page 589
- Section 5.9.7, “MSTP,” on page 781
- Section 5.10, “OpenFlow,” on page 814
- Section 5.12.3, “IGMP Snooping Querier,” on page 849
- the command “ip igmp snooping querier”
- the command “igmp snooping querier query-interval”

- the command “show ip igmp snooping querier”
- Section 5.13.2, “DCBX,” on page 871
- the command “lldp tlv-select dcbx”
- the command “dcb application-priority”
- the command “show dcb application-priority”
- Section 6.8, “DHCP Relay,” on page 1281

Updated:

- the command “show lldp interfaces”

Rev 1.7.0 – October 2013

Merged “MLNX-OS Command Reference Guide” Rev. 1.6.9 and “MLNX-OS User Manual” Rev. 1.6.9.

About this Manual

This manual provides general information concerning the scope and organization of this User's Manual.

Intended Audience

This manual is intended for network administrators who are responsible for configuring and managing Mellanox Technologies' SwitchX based Switch Platforms.

Related Documentation

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

Table 1 - Reference Documents

Document Name	Description
InfiniBand Architecture Specification, Vol. 1, Release 1.2.1	The InfiniBand Architecture Specification that is provided by IBTA.
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 SwitchX®-based switch system/series release note file
Mellanox Virtual Modular Switch Reference Guide	This reference architecture provides general information concerning Mellanox L2 and L3 Virtual Modular Switch (VMS) configuration and design.
Configuring Mellanox Hardware for VPI Operation Application Note	This manual provides information on basic configuration of the converged VPI networks.
MLNX-OS® XML API Reference Guide	This manual provides general information concerning MLNX-OS® XML API.

All of these documents can be found on the Mellanox website. They are available either through the product pages or through the support page with a login and password.

Glossary

Table 2 - Glossary

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.

Table 2 - Glossary

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	DCBX protocol is an extension of the Link Layer Discovery Protocol (LLDP). DCBX end points exchange request and acknowledgment messages. For flexibility, parameters are coded in a type-length-value (TLV) format.
DHCP	The Dynamic Host Configuration Protocol (DHCP) is an automatic configuration protocol used on IP networks.
Director Class Switch	A high density InfiniBand chassis switch system
DNS	Domain Name System. A hierarchical naming system for devices in a computer network
ETS	ETS provides a common management framework for assignment of bandwidth to traffic classes.
Fabric Management	The use of a set of tools (APIs) to configure, discover, and manage a group of devices organized as a connected fabric.
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.
Gateway	A network node that interfaces with another network using a different network protocol
GID (Global Identifier)	A 128-bit number used to identify a Port on a network adapter (see below), a port on a Router, or a Multicast Group.
GUID (Globally Unique Identifier)	A 64-bit number that uniquely identifies a device or component in a subnet
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
IB	InfiniBand
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).

Table 2 - Glossary

LDAP	The Lightweight Directory Access Protocol is an application protocol for reading and editing directories over an IP network.
LID (Local Identifier)	A 16 bit address assigned to end nodes by the subnet manager Each LID is unique within its subnet.
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
PFC/FC	Priority Based Flow Control applies pause functionality to traffic classes OR classes of service on the Ethernet link.
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
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.
SA (Subnet Administrator)	The interface for querying and manipulating subnet management data
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.
SM (Subnet Manager)	An entity that configures and manages the subnet, discovers the network topology, assign LIDs, determines the routing schemes and sets the routing tables. There is only one master SM and possible several slaves (Standby mode) at a given time. The SM administers switch routing tables thereby establishing paths through the fabric
SNMP	Simple Network Management Protocol. A network protocol for the management of a network and the monitoring of network devices and their functions

Table 2 - Glossary

NTP	Network Time Protocol. A protocol for synchronizing computer clocks in a network
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.
XML Gateway	Extensible Markup Language Gateway. Provides an XML request-response protocol for setting and retrieving HW management information.

1 Introduction

Mellanox® Operating System (MLNX-OS®) enables the management and configuration of Mellanox Technologies' SwitchX® Family silicon based switch platforms. MLNX-OS supports the Virtual Protocol Interconnect (VPI) technology which enables it to be used for both Ethernet and InfiniBand technology providing the user with greater flexibility.

MLNX-OS provides a full suite of management options, including support for Mellanox's Unified Fabric Manager® (UFM), SNMPv1, 2, 3, and web user interface (WebUI). In addition, it incorporates a familiar industry-standard CLI, which enables administrators to easily configure and manage the system.

1.1 System Features

Table 3 - General System Features

Feature	Description
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 controlsPower management - Auto-temperature control - High availability
Network Management Interfaces	- SNMP v1,v2c,v3 - interfaces (XML Gateway) - Puppet Agent
Security	- SSH - Telnet - RADIUS - TACACS+
Date and Time	NTP
Cables & Transceivers	Transceiver info
Unbreakable links	LLR
Virtual Port Interconnect® (VPI)	- Ethernet - InfiniBand

1.2 Ethernet Features

Table 4 - Ethernet Features

Feature	Description
General	• ACL – 6400 rules (permit/deny) • Breakout cables • Jumbo Frames (9K)
Ethernet support	• 48K unicast MAC addresses on SwitchX®-2 based systems • 2K static multicast MAC addresses • 90100 unicast MAC addresses on Spectrum™ based systems • DCBX • DHCP Relay • ETS (802.1Qaz) • Flow control (802.3x) • IGMP snooping v1,2 • LAG/LACP (802.3ad), 16 links per LAG (64 LAGs) • LLDP • MLAG • MSTP • OpenFlow 1.3 • PFC (802.1Qbb) • Rapid Spanning Tree (802.1w) • sFlow • VLAN (802.1Q) – 4K
IP routing	• BGP • DHCP Relay • ECMP • IGMP • IPv4 • IPv6 • OSPF • PIM • VLAN interface • Loopback interface • Router interface • VRRP

1.3 InfiniBand Features

Table 5 - InfiniBand Features

Feature	Description
Subnet Manager	• OpenSM • Partitions • High Availability Subnet Manager
Fabric diagnostics	• Fabric inspector

1.4

Figure 1: Managing an Ethernet Fabric Using MLNX-OS




Figure 2: Managing an InfiniBand Software Using MLNX-OS




Figure 3: Managing a

2 Getting Started

The procedures described in this chapter 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

➤ **To configure the switch:**

Step 1. Connect the host PC to the console (RJ-45) port of the switch system using the supplied cable. The console ports for systems are shown below.

Step 2.

Figure 4: Console Ports for CS75x0 Managed Systems




Figure 5: Console Ports for SB7700 Managed Systems





Figure 6: Console Ports for SX60xx/SX65xx Managed Systems



On SX65xx systems, when having dual management systems, first connect the cable and configure the master card and only then configure the slave. By default the master card is the top management module. Initial configuration must be performed on all of the management modules.



Make sure to connect to the console RJ-45 port of the switch and not to the MGT port.



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.

Step 3. Configure a serial terminal with the settings described below.



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

Table 6 - Serial Terminal Program Configuration for PPC Based Systems

Parameter	Setting
Baud Rate	9600
Data bits	8
Stop bits	1

Table 6 - Serial Terminal Program Configuration for PPC Based Systems

Parameter	Setting
Parity	None
Flow Control	None

Table 7 - Serial Terminal Program Configuration for x86 Based Systems

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

Step 4. You are prompted with the boot menu.

Mellanox MLNX-OS 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 “1” to boot with software version installed on partition #1.
 Select “2” to boot with software version installed on partition #2.
 Selecting “u” is not currently supported.
 Select “c” to proceed to advanced booting options – available to Mellanox Support only.

The MLNX-OS Boot Menu features a countdown timer. It is recommended to allow the timer to run out by not selecting any of the options.

Step 5. 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”.

Step 6. Go through the Mellanox configuration wizard.

The following table shows an example of a wizard session.

Table 8 - Configuration Wizard Session - IP Configuration by DHCP (Sheet 1 of 2)

Wizard Session Display (Example)	Comments
Mellanox configuration wizard 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 “y” and then press <Enter>.
Step1: 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 Table 9. If you enter “no” (no Zeroconf), then you need to enter a <i>static</i> IP, and the session will continue as shown in Table 10.
Step 3: Enable IPv6 [yes]	Perform this step to enable IPv6 on management ports. If you wish to enable IPv6, type “yes” and press <Enter>. If you enter “no” (no IPv6), then you will automatically be referred to Step 5.
Step 4: Enable IPv6 autoconfig (SLAAC) on mgmt0 interface	Perform this step to enable Stateless address autoconfig on external management port. If you wish to enable it, type “yes” and press <Enter>. If you wish to disable it, enter “no”.
Step 5: Use DHCPv6 on mgmt0 interface? [yes]	Perform this step to enable DHCPv6 on the MGMT0 interface.

Table 8 - Configuration Wizard Session - IP Configuration by DHCP (Sheet 2 of 2)

Wizard Session Display (Example)	Comments
Step 5: Admin password (Press <Enter> to leave unchanged)? <new_password> Step 4: Confirm admin password? <new_password>	To avoid illegal access to the machine, please type a password and then press <Enter>. Then confirm the password by re-entering it. Note that password characters are <i>not</i> printed.
You have entered the following information: 1. Hostname: <switch name> 2. Use DHCP on mgmt0 interface: yes 3. Enable IPv6: yes 4. Enable IPv6 autoconfig (SLAAC) on mgmt0 interface: yes 5. Enable DHCPv6 on mgmt0 interface: no 6. 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. Note: To run the command “configuration jump-start” you must be in Config mode.

Table 9 - Configuration Wizard Session - IP Zeroconf Configuration

Wizard Session Display - IP Zeroconf Configuration (Example)
Mellanox 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: Admin password (Enter to leave unchanged)? You have entered the following information: 1. Hostname: switch-112126 2. Use DHCP on mgmt0 interface: no 3. Use zeroconf on mgmt0 interface: yes 4. Default gateway: 192.168.10.1 5. Primary DNS server: 6. Domain name: 7. Enable IPv6: yes 8. Enable IPv6 autoconfig (SLAAC) on mgmt0 interface: yes 9. 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... <switch name> [standalone: master] >

Table 10 - Configuration Wizard Session - Static IP Configuration

Wizard Session Display - Static IP Configuration (Example)
Mellanox 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? [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: Admin password (Enter to leave unchanged)? You have entered the following information: 1. Hostname: switch-112126 2. Use DHCP on mgmt0 interface: no 3. Use zeroconf on mgmt0 interface: no 4. Primary IP address: 192.168.10.4 5. Netmask: 255.255.255.0 6. Default gateway: 192.168.10.1 7. Primary DNS server: 8. Domain name: 9. Enable IPv6: yes 10. Enable IPv6 autoconfig (SLAAC) on mgmt0 interface: no 11. 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... <switch name>[standalone: master] >

- Step 7.** 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 state
Admin up:          yes
Link up:           yes
IP address:        169.254.15.134
Netmask:           255.255.0.0
IPv6 enabled:      yes
Autoconf enabled:  yes
Autoconf route:    yes
Autoconf privacy:  no
IPv6 addresses:    1
IPv6 address:      fe80::202:c9ff:fe11:a1b2/64
Speed:             1000Mb/s (auto)
Duplex:            full (auto)
Interface type:    ethernet
Interface source:  physical
MTU:               1500
HW address:        00:02:C9:11:A1:B2
Comment:
RX bytes:          11700449      TX bytes:          15139846
RX packets:        55753        TX packets:        28452
RX mcast packets:  0           TX discards:       0
RX discards:       0           TX errors:          0
RX errors:         0           TX overruns:        0
RX overruns:       0           TX carrier:         0
RX frame:          0           TX collisions:      0
TX queue len:     1000
```

2.1.1 Re-Running the Wizard

- *To rerun the wizard:*

- Step 1.** Enter the config mode.

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

- Step 2.** Rerun the wizard.

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

2.2 Starting the Command Line (CLI)

- Step 1.** Set up an Ethernet connection between the switch and a local network machine using a standard RJ-45 connector.
- Step 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>
```

- Step 3.** Login to the switch (default username is *admin*, password *admin*)

- Step 4.** Read and accept the EULA when prompted.
- Step 5.** Once you get the prompt, you are ready to use the system.

```
Mellanox MLNX-OS Switch Management

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

Mellanox Switch
Please read and accept the Mellanox End User License Agreement located at:
http://www.mellanox.com/related-docs/prod_management_software/MLNX-OS_EULA.pdf

switch >
```

2.3 Starting the Web User Interface (WebUI)

- *To start a WebUI connection to the switch platform:*




WebUI access is enabled by default.

To disable web access, run the command “no web http enable” or “no web https enable” through the CLI.

- Step 1.** Set up an Ethernet connection between the switch and a local network machine using a standard RJ-45 connector.
- Step 2.** Open a web browser – Firefox 12, Chrome 18, IE 8, Safari 5 or higher.
Note: Make sure the screen resolution is set to 1024*768 or higher.
- Step 3.** Type in the IP address of the switch or its DNS name in the format: `http://<switch_IP_address>`.
- Step 4.** Login to the switch (default user name is *admin*, password *admin*).

Figure 7: MLNX-OS Login Window



The screenshot shows the Mellanox MLNX-OS Management Console login interface. At the top, the Mellanox logo is on the left, and the text "Mellanox MLNX-OS Management Console" is on the right. Below this, a status bar shows "Host: switch-Sea580" and "User: (not logged in)" with a "Login" button. A navigation bar contains icons for Setup, System, Security, Ports, Status, IB SM MGMT, Fabric Inspector, and ETH MGMT. The main content area is titled "Login" and contains the instruction "Please enter your username and password, then click 'Login'". Below this is a form with two input fields: "Account" with the value "admin" and "Password" which is empty. A "Login" button is positioned below the form. At the bottom of the page, there is a footer with the text "Mellanox MLNX-OS Switch Management" and "© 2009-2012 Mellanox Technologies, Inc." on the left, and "Best viewed using Firefox, Chrome, IE 7 or higher at 1024x768 resolution or higher." on the right.

Mellanox MLNX-OS Management Console

Host: switch-Sea580 User: (not logged in) Login

Setup System Security Ports Status IB SM MGMT Fabric Inspector ETH MGMT

Login

Please enter your username and password, then click "Login"

Account: admin

Password:


Login

Mellanox MLNX-OS Switch Management . © 2009-2012 Mellanox Technologies, Inc.

Best viewed using Firefox, Chrome, IE 7 or higher at 1024x768 resolution or higher.

- Step 5.** Read and accept the EULA if prompted.
You are only prompted if you have not accessed the switch via CLI before.

Figure 8: EULA Prompt



- Step 6.** The Welcome popup appears. After reading through the content, click OK to continue.
You may click on the links under Documentation to reach the MLNX-OS documentation.
The link under What's New takes you straight to the RN Changes and New Features section.


Figure 9: Welcome Popup



You may also tick the box to not show this popup again. But should you wish to see this window again, click “Product Documents” on the upper right corner of the WebUI.

Step 7. A default status summary is displayed as shown in Figure 10.

Figure 10: Display After Login



2.4 Licenses



Gateway is not supported in MLNX-OS® release 3.4.1110.

MLNX-OS software package can be extended with premium features. Installing a license allows you to access the specified premium features.



This section is relevant only to switch systems with an internal management capability.

The following licenses are offered with MLNX-OS software:

Table 11 - MLNX-OS Licenses

OPN	Valid on Product	Description
UPGR-6012-GW	SX6012	Ethernet L2/L3, Gateway
UPGR-1012-GW	SX1012	InfiniBand, Ethernet L3, Gateway
UPGR-6018-GW	SX6018	Ethernet L2/L3, Gateway
UPGR-6036-GW	SX6036	Ethernet L2/L3, Gateway
UPGR-1036-GW	SX1036	InfiniBand, Ethernet L3, Gateway
UPGR-1710-GW	SX1710	InfiniBand, Ethernet L3, Gateway
UPGR-6710-GW	SX6710	InfiniBand, Ethernet L3, Gateway
LIC-fabric-inspector	SX6036F/T; 6012F/T; 6018F/T; SB7700; SX65xx; CS75x0	InfiniBand fabric inspector monitoring and health
UPGR-xxxx-FCOE-J	All systems support- ing Ethernet directly or via license.	Enables FCoE protocol

2.4.1 Installing MLNX-OS® License (CLI)

➤ **To install an MLNX-OS license via CLI:**

Step 1. Login as *admin* and change to *Config* mode.

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

Step 2. Install the license using the key. Run:

```
switch (config) # license install <license key>
```


Step 3. Display the installed license(s) using the following command.

```
switch (config) # show licenses
License 1: <license key>
Feature: EFM_SX
Valid: yes
Active: yes
switch (config) #
```

Make sure that the “Valid” and “Active” fields both indicate “yes”.

Step 4. Save the configuration to complete the license installation. Run:

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



If you do not save the installation session, you will lose the license at the next system start up.


2.4.2 Installing MLNX-OS License (Web)

➤ **To install an MLNX-OS license via WebUI:**

Step 1. Log in as *admin*.

Step 2. Click the **Setup** tab and then **Licensing** on the left side navigation pane.

Figure 11: No Licenses Installed



Mellanox MLNX-OS Management Console

Host: [redacted] User: [redacted] Logout

Subnet Manager is running remotely.

Setup System Security Ports Status IB SM Mgmt Fabric Inspect ETH Mgmt IP Route Gateway Save

Licensing Product Documents

System Serial Number: [redacted]

Installed Licenses

License: No licenses installed.

[Remove](#)

[Add New License\(s\)](#)

Please enter one or more licenses, each on a separate line.

[Add Licenses](#)


© 2009-2016 Mellanox Technologies, Inc.

- Step 3.** Enter your license key(s) in the text box. If you have more than one license, please enter each license in a separate line. Click “Add Licenses” after entering the last license key to install them.



If you wish to add another license key in the future, you can simply enter it in the text box and click “Add Licenses” to install it.


Figure 12: Enter License Key(s) in Text Box



The screenshot shows the Mellanox MLNX-OS SX6506 Management Console interface. The top navigation bar includes the Mellanox logo, the title 'Mellanox MLNX-OS SX6506 Management Console', and user information: 'Host: switch-113dc8 User: admin Logout'. Below the navigation bar, there are tabs for 'Standalone', 'Virtual IP Active node', and 'Chassis master'. The main menu on the left includes 'Setup', 'System', 'Security', 'Ports', 'Status', 'IB SM MGMT', 'Fabric Inspector', and 'ETH MGMT'. The 'Licensing' section is active, showing a sidebar with various configuration options like 'Interfaces', 'HA', 'Routing', 'DNS', 'Hostname', 'Hosts', 'ARP', 'Neighbors', 'Virtual Switch Mgmt', 'Web', 'SNMP', 'Email Alerts', 'XML gateway', 'Logs', 'Configurations', 'Date and Time', 'NTP', and 'Licensing'. The main content area displays the 'System Serial Number' as 'M...' and the 'Installed Licenses' table. The table has columns for 'License', 'Key', 'Feature', 'Valid', 'Max num ufm ports supported', and 'Active'. The installed license is 'EFM_SX' with a key ending in '7U'. Below the table, there is a 'Remove' button and a section for 'Add New License(s)'. The 'Add New License(s)' section includes a text box with the placeholder '<your license key>' and an 'Add Licenses' button. A 'Save' button is located at the bottom right of the page.

All installed licenses should now be displayed.

Figure 13: Installed License



Step 4. Save the configuration to complete the license installation.



If you do not save the installation session, you will lose the installed licenses at the next system boot.

2.4.3 Retrieving a Lost License Key

In case of a lost MLNX-OS® license key, contact your authorized Mellanox reseller and provide the switch's *chassis serial number*.

➤ **To obtain the switch's chassis serial number:**

Step 1. Login to the switch.

Step 2. Retrieve the switch's *chassis serial number* using the command "show inventory".

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

Module	Part number	Serial Number	Asic Rev.	HW Rev.
CHASSIS	MSX1036B-1SFR	MT1205X01549	N/A	A1
MGMT	MSX1036B-1SFR	MT1205X01549	0	A1
FAN	MSX60-FF	MT1206X07209	N/A	A3
PS1	MSX60-PF	MT1206X06697	N/A	A2

```
switch (config) #
```

Step 3. Send your Mellanox reseller the following information to obtain the license key:

- The chassis serial number

- The type of license you need to retrieve. Refer to [“Licenses” on page 63](#).

Step 4. Once you receive the license key, you can install the license as described in the sections above.

2.4.4 Commands

file eula upload

file eula upload <filename> <URL>

Uploads the Mellanox End User License Agreement to a specified remote location.

Syntax Description	filename	The Mellanox End User License Agreement
	URL	URL or scp://username[:password]@hostname/path/filename
Default	N/A	
Configuration Mode	Config	
History	3.4.1100	
Role	monitor/admin	
Example	<pre>switch (config) # file help-docs upload Mellanox_End_User_ License_Agreement.pdf <scp://username[:password]@hostname/path/ filename> switch (config) #</pre>	
Related Commands	license	
Note		

file help-docs upload

file help-docs upload <filename> <URL or scp://username[:password]@hostname/path/filename>

Uploads the MLNX-OS UM or RN 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	
Role	admin	
Example	<pre>switch (config) # file help-docs upload MLNX-OS_VPI_User_Manual.pdf <scp://username[:password]@hostname/path/filename> switch (config) #</pre>	
Related Commands		
Note		

license delete

license delete <license-number>

Removes license keys by ID.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.4.1100
Role	admin
Example	<pre>switch (config) # license delete <license-key> switch (config) #</pre>
Related Commands	
Note	Before deleting a license from a switch which is configured to a system profile other than its default, the user must first disable all interfaces and then return the switch to its default system profile.

license install

license install <license-key>

Installs a new license key.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.4.1100
Role	admin
Example	switch (config) # licenses install <license-key> switch (config) #
Related Commands	
Note	

show licenses

show licenses

Displays a list of all installed licenses. For each license, the following is displayed:

- a unique ID which is a small integer
- the text of the license key as it was added
- whether or not it is valid and active
- which feature(s) it is activating
- a list of all licensable features specifying whether or not it is currently activated by a license

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.4.1100
Role	admin
Example	<pre>switch (config) # show licenses License 1: <license key> Feature: SX_CONFIG Valid: yes Active: yes switch (config) #</pre>
Related Commands	
Note	

3 User Interfaces

3.1 Command Line Interface Overview

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

3.1.1 CLI Modes

The CLI can be in one of following modes, and each mode makes available a certain group (or level) of commands for execution. The different CLI configuration modes are:

Table 12 - CLI Modes and Config Context

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 <code>enable</code> 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 <code>configure terminal</code> command moves the user from Enable mode to Config mode. Config mode is allowed only for user accounts in 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 <code>exit</code> or <code>no configure</code> . Note that moving directly from/to Standard mode to/from Config mode is not possible.
Config Interface Management	Configuration mode for management interface <code>mgmt0</code> , <code>mgmt1</code> and <code>loopback</code> .
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” can be applied within any context.

3.1.2 Syntax Conventions

To help you identify the parts of a CLI command, this section explains conventions of presenting the syntax of commands.

Table 13 - Syntax Conventions

Syntax Convention	Description	Example
< > Angled brackets	Indicate a value/variable that must be replaced.	<1...65535> or <switch interface>
[] Square brackets	Enclose optional parameters. However, only one parameter out of the list of parameters listed can be used. The user cannot have a combination of the parameters unless stated otherwise.	[destination-ip destination-port destination-mac]
{ } Braces	Enclose 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 type the angled or square brackets, vertical bar, or braces in command lines. This guide uses these symbols only to show the types of entries.



CLI commands and options are in lowercase and are case-sensitive. For example, when you enter the **enable** command, enter it all in lowercase. It cannot be **ENABLE** or **Enable**. Text entries you create are also case-sensitive.

3.1.3 Getting Help

You may request context-sensitive help at any time by pressing “?” on the command line. This will show a list of choices for the word you are on, or a list of top-level commands if you have not typed anything yet.

For example, if you are in Standard mode and you type “?” at the command line, then you will get the following list of 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-11a596 [standalone: master] >
```

If you type a legal string and then press “?” *without* a space character before it, then you will either get a description of the command that you have typed so far or the possible command/parameter completions. If you press “?” *after* a space character and “<cr>” is shown, this means that what you have entered so far is a complete command, and that you may press Enter (carriage return) to execute it.

Try the following to get started:

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

You can also 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 (more than one completion option is available, that is), then you can hit double tabs to obtain the disambiguation options. Thus, if you are in Enable mode and wish 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
```

(There are three commands that start with the letter “c”: clear, cli and configure.)

3.1.4 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 you press <Enter>.

If an error is encountered in executing a command, the response will begin with “%”, followed by some text describing the error.

3.1.5 Using the “no” Form

Several Config mode commands offer the negation form using the keyword “no”. This no form can be used to disable a function, to cancel certain command parameters or options, or to reset a parameter value to its default. To re-enable a function or to set cancelled command parameters or options, enter the command without the “no” keyword (with parameter values if necessary).

The following example performs the following:

1. Displays the current CLI session options.
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-1 [standalone: master] (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
...

```

3.1.6 Parameter Key

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

Table 14 - Angled Brackets Parameter Description

Parameter	Description
<domain>	A domain name, e.g. “mellanox.com”.
<hostname>	A hostname, e.g. “switch-1”.
<ifname>	An interface name, e.g. “mgmt0”, “mgmt1”, “lo” (loopback), etc.
<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: “debug”, “info”, “notice”, “warning”, “error”, “crit”, “alert”, “emerg”.
<GUID>	Globally Unique Identifier. A number that uniquely identifies a device or component.

Table 14 - Angled Brackets Parameter Description

Parameter	Description
<MAC address>	A MAC address. The segments may be 8 bits or 16 bits at a time, and may be delimited by “:” or “.”. So you could say “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”). These two 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 man page. (The value you provide 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. The implementation of ftp does not support authentication, so use scp or sftp for that. Note also that if you omit the “:password” part, you may be prompted for the password in a follow up prompt, where you can type it 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 you do not want any password. If you include the “:” character, this will be taken as an explicit declaration that the password is empty, and you will not be prompted in any case.

3.1.7 CLI Pipeline Operator Commands

3.1.7.1 “include” and “exclude” CLI Filtration Options

The MLNX-OS 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 (i.e. “|”). Quotation marks may be used to include or exclude a string including space, and multiple filters can be used simultaneously. For example:

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

Examples:

```
switch (config) # show asic-version | include SX
MGMT          SX          9.3.3150

arc-switch14 [standalone: master] (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
--
Tx
0 packets
0 unicast packets
0 multicast packets
0 bytes
```

3.1.7.2 “watch” CLI Monitoring Option

MLNX-OS also allows viewing a live feed of the progress of any “show” command by using the “watch” option as follows:

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

Running the command as such displays an output of the show command that gets updated at a time interval which may be specified using the “interval” parameter (2 seconds by 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 Capacity Feed Status
      [Watts] [Volts] [Amp]  [Watts]
-----
PS1    power-mon       input 85.00 230.00 0.38 460.00 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: 462.00 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 also be used in conjunction with the “include” and “exclude” options as follows:

```
switch (config) # <show command> | {include | exclude} <extended regular expression> |
watch [diff] [interval <1-100 secs>]
```

For example:

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

3.1.7.3 “json-print” CLI Option



This feature is available on x86 based systems only.

The MLNX-OS 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. For 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 [Section 4.18.2, “JSON API,” on page 558](#).

For a list of commands supporting the JSON API, please refer to [Appendix C, “Show Commands Supported by JSON API,”](#) on page 1646.

3.1.8 CLI Shortcuts

Table 15 presents the available keyboard shortcuts on the MLNX-OS® CLI.

Table 15 - CLI Keyboard Shortcuts

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 auto-complete 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-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
Esc [A	Next line (same as DOWN ARROW)
Esc [B	Next line (same as UP ARROW)

Table 15 - CLI Keyboard Shortcuts

Key Combination	Description
Esc [C	Move forward one character from cursor
Esc [D	Move backward one character from cursor

3.2 Web Interface Overview

MLNX-OS® package equipped with web interface which is a web GUI that accept input and provide output by generating webpages which can be viewed by the user using a web browser.

The following web browsers are supported:

- Internet Explorer 8.0 or higher
- Chrome 18 or higher
- Mozilla Firefox 12 or higher
- Safari 5 or higher


The web interface makes available the following perspective tabs:

- Setup
- System
- Security
- Ports
- Status
- IB SM Management
- Fabric Inspector
- Ethernet Management
- IP Route
- Gateway



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

Figure 14: WebUI



3.2.1 Setup Menu

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

Table 16 - WebUI Setup Submenus

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.
HA	Creates, joins or modifies an InfiniBand subnet.
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.
DNS	Configures, removes, modifies or displays static and dynamic name servers.

Table 16 - WebUI Setup Submenus

Submenu Title	Description
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.
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.
Logs	Sets up system log files, remote log sinks, and log formats.
Configurations	Manages, activates, saves, and imports MLNX-OS SwitchX configuration files, and executes CLI commands.
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 MLNX-OS licenses.

3.2.2 System Menu

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

Table 17 - WebUI System Submenus

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.

Table 17 - WebUI System Submenus

Submenu Title	Description
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.
MLNX-OS Upgrade	Displays the installed MLNX-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.

3.2.3 Security Menu

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

Table 18 - WebUI Security Submenus

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.

3.2.4 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.

Table 19 - WebUI Ports Submenus

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.

Table 19 - WebUI Ports Submenus

Submenu Title	Description
Monitor Session	Displays monitor session summary and enables configuration of a selected session.
Protocol Type	Manages the link protocol type

3.2.5 Status Menu

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

Table 20 - WebUI Status Submenus

Submenu Title	Description
Summary	Displays general information about the switch system and the MLNX-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.
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.
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.

3.2.6 IB SM Mgmt



The IB SM MGMT menu is not supported in Ethernet systems.

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

Table 21 - WebUI IB SM Mgmt Submenus

Submenu Title	Description
Summary	Displays the local Subnet Manager (SM) status (running time, failures, etc).
Base SM	Manages basic SM configuration (enabling SM, priority level, and restoring initial configuration).
Advanced SM	Manages basic SM configuration (enabling SM, priority level, and restoring initial configuration).
Expert SM	Configures security and GUID based prefixes (m_key, sm_key, sa_key, etc), and manages special SM attributes that should not be changed except by expert users of the Subnet Manager who understand the risks of manipulating these attributes.
Compute nodes	Adds compute nodes using network adapter port GUIDs.
Root nodes	Adds root nodes using switch GUIDs.
Partitions	Manages partition keys (sets removes or displays the partition keys).
Basic Qos	Configures basic QoS attributes such as default QoS settings, and VL arbitration low and high entries. It also displays and manages SL-to-VL mappings.

3.2.7 Fabric Inspector



The Fabric Inspctr menu is not applicable when the switch profile is not InfiniBand or VPI.



The Fabric Inspctr menu requires a license (LIC-fabric-inspector).

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

Table 22 - WebUI Fabric Inspctr Submenus

Submenu Title	Description
Summary	Displays a fabric status summary, including the time of last fabric update, what systems are in the fabric, what InfiniBand devices are identified, etc.
IB Systems	Displays information about all identified InfiniBand systems in the fabric (adapters, switches, etc).
IB Nodes	Displays information about InfiniBand nodes in the fabric. It is possible to filter display by the type of InfiniBand node (HCA adapter, switch, etc).
IB Ports	Displays all active InfiniBand ports in the fabric. It is possible to filter display by the type of InfiniBand port (HCA port, switch port, switch management port, etc), by the port rate (speed or width), by the Subnet Manager status on the node, by node traffic, etc.
Connections	Displays all active connections in the fabric. It is possible to filter display by the link type (switch to switch, switch to HCA, etc) and by the link rate (speed or width).
System Names	Allows the mapping of System Names to GUIDs to ease system identification.

3.2.8 ETH Mgmt



The Eth Mgmt menu is not applicable when the switch profile is not Ethernet or VPI.

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

Table 23 - WebUI ETH Mgmt Submenus

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.
IGMP Snooping	Manages IGMP snooping in the switch.
ACL	Manages Access Control in the switch.

Table 23 - WebUI ETH Mgmt Submenus

Submenu Title	Description
Priority Flow Control	Manages priority flow control.

3.2.9 IP Route

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

Table 24 - WebUI IP Route Submenus

Submenu Title	Description
Router Global	Enables/disables IP Routing protocol on the machine.
IP Route	Not implemented.
IP Interface	Not implemented.
Address Resolution	Not implemented.
IP Diagnostic	Not implemented.

3.3 Secure Shell (SSH)



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

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:*

Step 1. Change to Config mode Run:

```
switch [standalone: master] > enable
switch [standalone: master] # configure terminal
switch [standalone: master] (config) #
```

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

```
switch [standalone: master] (config) # ssh client global known-host "myserver ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAIEAsXeklqc8T0EN2mnMcVcfhueaRYzIVqt4rVsrERIjmlJh4mkYYIa8hGGikN
a+t5xw2dRrNxnHYLK51bUsSG1ZNwZT1Dpme3pAZeMY7G4ZMgGIW9xOuaXgAA3eBeoUjFdi6+1BqchWk0nTb+gm
fI/MK/heQNns7AtTrvqg/05ryIc="
switch [standalone: master] (config) #
```

Step 3. Verify what keys exist in the host. Run:

```
switch [standalone: master] (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.

switch [standalone: master] (config) #
```

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:*

Step 1. Connect to the system from the host SSH.

Step 2. Add the -h parameter after the cli (as shown in the example below) 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 Management Information Bases (MIBs)

The inventory in the switch system can be accessed through a MIB browser. These devices are indexed (entPhysicalIndex) using three levels:

1. Module layer which includes modules located on system (e.g. cables, fan, power supply, etc.). See table [Table 25](#) for more details.
2. Device layer which includes system devices (e.g. switch devices, sensor aggregators, etc.). See table [Table 26](#) for more details.
3. Sensor layer which includes system sensors (e.g. fan, and temperature sensors) located in the devices. See table [Table 27](#) for more details.

Each layer is assigned a fixed position in the index number to represent it.

Figure 15: Index Scheme

Mod. Type	2-Digit Module Index		Device Name		Device Index #1	Device Index #2	Sensor Type	Sensor Index
1	2	3	4	5	6	7	8	9

Each position could indicate different types of component according to the following criteria:

Table 25 - Module Type

Number	Description
1	Chassis
2	Management
3	Spine
4	Leaf
5	Fan
6	Power supply
7	BBU
8	x86 CPU
9	Port module

Table 26 - Device Type

Number	Description
01	PS
02	FAN
03	BOARD_MONITOR
04	CPU_BOARD_MONITOR
05	SX
06	SIB
07	CPU_MEZZ_TEMP
08	CPU Package Sensor
09	CPU Core Sensor
10	SX_AMBIENT_TEMP
11	SX_MONITOR
12	AUX_IN_TMP_SNSR
13	AUX_OUT_TMP_SNSR
14	MAIN_IN_TMP_SNSR
15	MAIN_OUT_TMP_SNSR
16	CPU_MEZZ_TEMP
17	Controller
18	QSFP_TEMP
19	QSFP-ASIC

Table 26 - Device Type

Number	Description
20	Board AMB temp
21	Ports AMB temp
22	Power monitor
23	PS_MONITOR
24	SWB AMB temp
25	pcie-switch-temp
26	SPC

Table 27 - Sensor Type

Number	Description
1	t – temperature sensor
2	f – fan sensor

For example:

- 401191311

The first layer is “401” where:

- “4”, according to [Table 25](#), indicates a leaf
- “01” indicates index #1 (Leaf #1)

The second layer is “1913” where:

- “19”, according to [Table 26](#), indicates a QSFP ASIC
- “1” indicates ASIC #1
- “3” indicates sensor #3 (QSFP-ASIC1-3)

The third layer is “11” where:

- “1”, according to [Table 27](#), indicates a temperature sensor
- “1” indicates sensor #1 (T1)

The resulting output in the entPhysicalDescr column of the MIB would be: L01/QSFP-ASIC-1/T1.

- 501020021

The first layer is 501 where

- “5”, according to [Table 25](#), indicates a fan
- “01 indicates index #1 (Fan #1)

The second layer is 0200 where:

- 02, according to [Table 26](#), indicates a fan
- 0 – indicates that there is no first index

- 0 – indicates that there is no second index

The third layer is 21 where:

- “2”, according to [Table 27](#), indicates a fan sensor
- “1” indicates sensor #1 (F1)

The resulting output in the entPhysicalDescr column of the MIB would be: FAN1/FAN/F1.

3.5 Commands

3.5.1 CLI Session

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

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
Role	admin
Example	switch (config) # cli clear-history switch (config) #
Related Commands	N/A
Note	

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 all future sessions.

The no form of the command deletes or disables the default CLI options.

Syntax	Description
minutes	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 and all future sessions “prefix-modes show-config” uses prefix modes in “show configuration” output for current and all future sessions
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
Role	admin

Example

```
switch (config) # cli default prefix-modes enable
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 future sessions:
  Auto-logout:            disabled
  Paging:                 enabled
  Progress tracking:      enabled
  Prefix modes:           enabled (and use in 'show configuration')

Settings for both this session and future ones:
  Show hidden config:     yes
  Confirm losing changes: yes
  Confirm reboot/shutdown: no
  Confirm factory reset:  yes
  Prompt on empty password: yes
switch (config) #
```

Related Commands `show cli`

Note

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: 3-60
Default	50 sessions	
Configuration Mode	Config	
History	3.5.0200	
Role	admin	
Example	switch (config) # cli max-sessions 40 switch (config) #	
Related Commands	show terminal	
Note		

cli session

```
cli session {auto-logout <minutes> | paging enable | prefix-modes {enable | show-  
config} | progress enable | terminal {length <size> | resize | type <terminal-type>  
| width} | x-display full <display>}  
no cli session {auto-logout | paging enable | prefix-modes {enable | show-config} |  
progress enable | terminal type | x-display}
```

Configures default CLI options for all future sessions.
The no form of the command deletes or disables the CLI sessions.

Syntax Description	minutes	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 and all future sessions “prefix-modes show-config” uses prefix modes in “show configuration” output for current and all future sessions
	progress enable	Enables progress updates.
	terminal length	Sets the number of lines for the current terminal. Valid range is 5-999.
	terminal resize	Resizes the CLI terminal settings (to match the actual terminal window).
	terminal-type	Sets the terminal type. Valid options are: - ansi - console - dumb - linux - unknown - vt52 - vt100 - vt102 - vt220 - vt320 - xterm
	terminal width	Sets the width of the terminal in characters. Valid range is 34-999.
	x-display full <display>	Specifies the display as a raw string, e.g local-host:0.0.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # cli session auto-logout switch (config) #</pre>	

Related Commands show terminal

Note

terminal

terminal {length <number of lines> | resize | type <terminal type> | width <number of characters>}
no terminal type

Configures default CLI options for all future sessions.
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	
Role	admin	
Example	switch (config) # terminal length 500 switch (config) #	
Related Commands	show terminal	
Note		

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	Enabled
Configuration Mode	Config
History	3.4.3000
Role	admin
Example	switch (config) # terminal sysrq enable switch (config) #
Related Commands	show terminal
Note	

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
Role	admin
Example	<pre>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 future sessions: Auto-logout: disabled Paging: enabled Progress tracking: enabled Prefix modes: enabled (and use in 'show configuration') Settings for both this session and future ones: Show hidden config: yes Confirm losing changes: yes Confirm reboot/shutdown: no Confirm factory reset: yes Prompt on empty password: yes switch (config) #</pre>
Related Commands	cli default
Note	

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
Role	admin
Example	switch (config) # show cli max-sessions Maximum number of CLI sessions: 50 switch (config) #
Related Commands	
Note	

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
Role	admin
Example	switch (config) # show cli num-sessions Current number of CLI sessions: 40 switch (config) #
Related Commands	
Note	

3.5.2 Banner

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	string	Text string.
Default	“Mellanox MLNX-OS Switch Management”	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # banner login Example switch (config) #	
Related Commands	show banner	
Note	If more than one word is used (there is a space) quotation marks should be added (i.e. “xxxx xxxx”).	

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.

Syntax Description	string	Text string.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
	3.5.0200	Added no form of the command
Role	admin	
Example	switch (config) # banner login-local Testing switch (config) #	
Related Commands	show banner	
Note	- 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").	

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.

Syntax Description	string	Text string.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
	3.5.0200	Added no form of the command
Role	admin	
Example	switch (config) # banner login-remote Testing switch (config) #	
Related Commands	show banner	
Note	- 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").	

banner logout

banner logout <string>
no banner logout

Set system logout banner (for both local and remote logins).
The no form of the command resets the banner.

Syntax Description	string	Text string.
Default	N/A	
Configuration Mode	Config	
History	3.5.0200	
Role	admin	
Example	switch (config) # banner logout Testing switch (config) #	
Related Commands	show banner	
Note	If more than one word is used (there is a space) quotation marks should be added (i.e. "xxxx xxxx").	

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.

Syntax Description	string	Text string.
Default	N/A	
Configuration Mode	Config	
History	3.5.0200	
Role	admin	
Example	switch (config) # banner logout-local Testing switch (config) #	
Related Commands	show banner	
Note	• 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").	

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.

Syntax Description	string	Text string.
Default	N/A	
Configuration Mode	Config	
History	3.5.0200	
Role	admin	
Example	switch (config) # banner logout-remote Testing switch (config) #	
Related Commands	show banner	
Note	- 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").	

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	"Mellanox Switch"	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # banner motd "My Banner"	
Related Commands	show banner	
Note	- 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.	

show banner

show banner

Displays configured banners.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000 3.5.0200 Updated Example
Role	Any Command Mode
Example	<pre>switch (config) # show banner Banners: Message of the Day (MOTD): Mellanox Switch Login: Mellanox MLNX-OS Switch Management Logout: Goodbye switch (config) #</pre>
Related Commands	<pre>banner login banner login-local banner login-remote banner logout banner logout-local banner logout-remote banner motd</pre>
Note	

3.5.3 SSH

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
Role	admin
Example	<pre>switch (config) # ssh server enable 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 server ports: 22 Interface listen enabled: yes No Listen Interfaces. Host Key Finger Prints: RSA v1 host key: SHA256:ElFoK7Jts7ejIws0Jgs3yt46go0Ckln0JzNzAGx0ue4 (2048) RSA v2 host key: SHA256:N4n+Un/lErjtzmmDJH+qcdsmHgHc0itLYArFgqP+UFI (2048) DSA v2 host key: SHA256:2rIuzmPD90AWooQaEjI1SH5EF0DjQ9DDSTaAMrzDFCY (1024) switch (config) #</pre>
Related Commands	show ssh server
Note	Disabling SSH server does not terminate existing SSH sessions, it only prevents new ones from being established.

ssh server host-key

ssh server host-key {<key-type> {private-key <private-key>| public-key <public-key>} | generate}

Manipulates host keys for SSH.

Syntax Description	key-type	- rsa1 - RSAv1 - 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
Role	admin	

Example

```
switch (config) # ssh server host-key dsa2 private-key
Key: *****
Confirm: *****
switch (config) # show ssh server host-keys
SSH server configuration:
  SSH server enabled:      yes
  Minimum protocol version: 2
  X11 forwarding enabled:  no
  SSH server ports:       22

Interface listen enabled: yes
No Listen Interfaces.

Host Key Finger Prints:
  RSA v1 host key: a0:63:db:96:e2:95:5a:5a:fd:a8:d0:f4:ab:e3:5f:f8
  RSA v2 host key: 1e:b7:8b:ec:ab:35:98:be:6b:d6:12:c2:18:72:12:d6
  DSA v2 host key: 7c:4a:f7:72:51:67:b5:0b:cd:a2:d2:b9:f3:be:3e:68

Host Keys:
  RSA v1 host key: "switch-5ea5d8 1024 35
12457497995374010105491416867919987976776882016984375942831915584962796
99375406596085804272219042450456598705866658144854493132172365068789517
13570509420864336951833046700451354269467758379288848962624165330724512
16091899983038691571036219385577978596282214644533444813712105628654158
3022982220576029771297093"
  RSA v2 host key: "switch-5ea5d8 ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAA-
IEArB9i50nukAHNUOkwpCmEl0m88kJgBzL22+F5tfaSn+S0pVYxrceZeyuzXsoZ1VtFTk2-
FydwY0YvMS0Kcv2PuCrPZV/
GYd31QEnn22rEmrLPrKCrMl1Xluy6DFlr30gwwm1baobmDLg/gSziWz/gc4Jgqf2CyX-
Fq4pzaR1jar1Vk="
  DSA v2 host key: "switch-5ea5d8 ssh-dss AAAAB3NzaC1kc3MAAAC-
BAMEJ3S+nyaHhRbwv3tJqLWttDC35RZVC5iG4ZEVMHhp28VL940cyyuGh39VCdM9pEVaI7h
zZrsgHrNqakb/YLD/7anGH3wpl9Fx8lfe0RH3bloJzG+mJ6R5momdoPCrKwEKiKABKE00-
jLzLVznpP0IHxjwF+TbR3dK5HwVzQYw/bAAAAFQCBoDPqBZZa+2KylKLzUsbZ2pKhgQAAA-
IAJK+StiQdt0Rw1B5UCMzTrTef5L07DSfVreMEYtTRnBBtgVSNqQFwpSQIYbVDHQR9T6qCM
4V039DuHUGQ1TMDIX7t+9mfbB87YyUu5a/ndbf3GhNhxHwWbzlr9hgLL7FSHA7DYH7bV0Z-
RlqxH64eQKGZqy1ps/F4E31lyn7GC4EQAAAIA/2osHipXf+NRjplgfmHR0VVf/mGE9Vzc9/
AMUxLJn5VhvEJ5CZW9cI+LxM0Jojh0j3YW3B1czGxR0bDA9vUbKXTNc8bkgoUrxySAH1rH
N0PqJgeT4L009AItSp3m1mxHqdS7jixfTv0TEKWxrgpczlmTB8+zjhUah/YuuB12H
g=="
switch (config) #
```

Related Commands

```
show ssh server
system secure-mode enable
```

Note

When working in secure mode, the commands “ssh server host-key rsa1” and “ssh server host-key generate” do not create RSAv1 key-type.

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 <inf>	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	
Role	admin	
Example	<pre>switch (config) # ssh server listen enable switch (config) # show ssh server SSH server configuration: SSH server enabled: yes Minimum protocol version: 2 X11 forwarding enabled: no SSH server ports: 22 Interface listen enabled: yes No Listen Interfaces. Host Key Finger Prints: RSA v1 host key: a0:63:db:96:e2:95:5a:5a:fd:a8:d0:f4:ab:e3:5f:f8 RSA v2 host key: 1e:b7:8b:ec:ab:35:98:be:6b:d6:12:c2:18:72:12:d6 DSA v2 host key: 7c:4a:f7:72:51:67:b5:0b:cd:a2:d2:b9:f3:be:3e:68 switch (config) #</pre>	
Related Commands	show ssh server	
Note		

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.
Default	6 attempts	
Configuration Mode	Config	
History	3.5.0200	
	3.5.1000	Increased minimum number of attempts allowed
Role	admin	
Example	switch (config) # ssh server login attempts 5	
Related Commands	show ssh server	
Note		

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	
Role	admin	
Example	switch (config) # ssh server login timeout 130	
Related Commands	show ssh server	
Note		

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	
Role	admin	
Example	<pre>switch (config) # ssh server min-version 2 switch (config) # show ssh server SSH server configuration: SSH server enabled: yes Minimum protocol version: 2 X11 forwarding enabled: no SSH server ports: 22 Interface listen enabled: yes No Listen Interfaces. Host Key Finger Prints: RSA v1 host key: a0:63:db:96:e2:95:5a:5a:fd:a8:d0:f4:ab:e3:5f:f8 RSA v2 host key: 1e:b7:8b:ec:ab:35:98:be:6b:d6:12:c2:18:72:12:d6 DSA v2 host key: 7c:4a:f7:72:51:67:b5:0b:cd:a2:d2:b9:f3:be:3e:68 switch (config) #</pre>	
Related Commands	show ssh server	
Note		

ssh server ports

ssh server ports {<port1> [<port2>...]}

Specifies which ports the SSH server listens on.

Syntax Description	port	Port number in [1...65535].
Default	22	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ssh server ports 22 switch (config) # show ssh server SSH server configuration: SSH server enabled: yes Minimum protocol version: 2 X11 forwarding enabled: no SSH server ports: 22 Interface listen enabled: yes No Listen Interfaces. Host Key Finger Prints: RSA v1 host key: a0:63:db:96:e2:95:5a:5a:fd:a8:d0:f4:ab:e3:5f:f8 RSA v2 host key: 1e:b7:8b:ec:ab:35:98:be:6b:d6:12:c2:18:72:12:d6 DSA v2 host key: 7c:4a:f7:72:51:67:b5:0b:cd:a2:d2:b9:f3:be:3e:68 switch (config) #</pre>	
Related Commands	show ssh server	
Note	- Multiple ports can be specified by repeating the <port> parameter - The command will remove any previous ports if not listed in the command	

ssh server security strict

ssh server security strict

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
Role	admin
Example	switch (config) # ssh server security strict switch (config) #
Related Commands	show ssh server
Note	The following ciphers are disabled for SSH when strict security is enabled: • aes256-cbc • aes192-cbc • aes128-cbc • arcfour • blowfish-cbc • cast128-cbc • rijndael-cbc@lysator.liu.se

ssh server tcp-forwarding enable

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
Role	admin
Example	switch (config) # ssh server tcp-forwarding enable switch (config) #
Related Commands	show ssh server
Note	

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	X11-forwarding is disabled.
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # ssh server x11-forwarding enable switch (config) # show ssh server SSH server configuration: SSH server enabled: yes Minimum protocol version: 2 X11 forwarding enabled: yes SSH server ports: 22 Interface listen enabled: yes No Listen Interfaces. Host Key Finger Prints: RSA v1 host key: a0:63:db:96:e2:95:5a:5a:fd:a8:d0:f4:ab:e3:5f:f8 RSA v2 host key: 1e:b7:8b:ec:ab:35:98:be:6b:d6:12:c2:18:72:12:d6 DSA v2 host key: 7c:4a:f7:72:51:67:b5:0b:cd:a2:d2:b9:f3:be:3e:68 switch (config) #</pre>
Related Commands	N/A
Note	

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	
Role	admin	

Example

```
switch (config) # ssh client global host-key-check no
switch (config) # ssh client global known-host "72.30.2.2 ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAIEArB9i50nukAHNUOkwpCmEl0m88kJgB-
zL22+F5tfaSn+S0pVYxrceZeyuzXsoZ1VtFTk2Fydwy0YvMS0Kcv2PuCrPZV/
GYd31QEnn22rEmrLPrKCrMl1XlUy6DFlr30gwwm1baobmDLG/gSziWz/gc4Jgqf2CyX-
Fq4pzaR1jar1Vk="

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.

switch (config) #
```

Related Commands

```
show ssh client
```

Note

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	
Role	admin	
Example	<pre>switch (config) # ssh client user admin known-host 172.30.1.116 remove switch (config) #</pre>	

Related Commands	show ssh client

Note	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.

slogin

slogin [<slogin options>] <hostname>

Invokes the SSH client. The user is returned to the CLI when SSH finishes.

Syntax Description	slogin options	usage: slogin [-1246AaCfGkNnqsTtVvXxY] [-b bind_address] [-c cipher_spec] [-D port] [-e escape_char] [-F configfile] [-i identity_file] [-L port:host:hostport] [-l login_name] [-m mac_spec] [-o option] [-p port] [-R port:host:hostport] [user@]hostname [command]
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	monitor/admin	
Example	<pre>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 Warning: Permanently added '192.168.10.70' (RSA) to the list of known hosts. Mellanox MLNX-OS Switch Management Last login: Sat Feb 28 22:55:17 2009 from 10.208.0.121 Mellanox Switch switch (config) #</pre>	
Related Commands	N/A	
Note		

show ssh client

show ssh client

Displays the client configuration of the SSH server.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>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. switch (config) #</pre>
Related Commands	N/A
Note	

show ssh server

show ssh server

Displays SSH server configuration.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000 3.4.0000 Updated Example 3.5.0200 Added SSH login timeout and max attempts
Role	admin
Example	<pre>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 ports: 22 Interface listen enabled: yes No Listen Interfaces. Host Key Finger Prints and Key Lengths: RSA v1 host key: 5f:4e:5f:4a:81:bb:6a:b4:06:52:77:eb:d3:ad:78:92 (2048) RSA v2 host key: 15:e2:a8:45:1c:58:1b:00:cc:29:ec:00:38:83:49:00 (2048) DSA v2 host key: df:c0:ac:a6:3e:a5:52:a5:d1:f6:22:37:ef:f1:08:f9 (1024) switch (config) #</pre>
Related Commands	ssh server
Note	

3.5.4 Remote Login

telnet

telnet

Logs into another system using telnet.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # (config) # telnet telnet>
Related Commands	telnet-server
Note	

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
Role	admin
Example	switch (config) # telnet-server enable switch (config) # show telnet-server Telnet server enabled: yes
Related Commands	show telnet-server
Note	

show telnet-server

show telnet-server

Displays telnet server settings.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show telnet-server Telnet server enabled: yes switch (config) #
Related Commands	telnet-server enable
Note	

3.5.5 Web Interface

web auto-logout

web auto-logout <number of minutes>
no web auto-logout <number of minutes>

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	number of minutes	The length of user inactivity in minutes. 0 will disable the inactivity timer (same as a “no web auto-logout” command).
Default	60 minutes	
Configuration Mode	Config	
History	3.1.0000 3.4.0000	Updated Example
Role	admin	
Example	<pre>switch (config) # web auto-logout 60 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: all HTTPS certificate name: default-cert Listen enabled: yes No Listen Interfaces. 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 switch (config) #</pre>	

Related Commands	show web

Note	The no form of the command does not automatically log users out due to inactivity.

web cache-enable

web cache-enable
no web cache-enable

Enables web clients to cache webpages.
The no form of the command disables web clients from caching webpages.

Syntax Description	N/A
Default	Enabled
Configuration Mode	Config
History	3.4.1100
Role	admin
Example	switch (config) # no web cache-enable
Related Commands	N/A
Note	

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
Role	admin
Example	switch (config) # web client cert-verify
Related Commands	N/A
Note	

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	
Role	admin	
Example	switch (config) # web client ca-list default-ca-list	
Related Commands	N/A	
Note		

web enable

web enable
no web enable

Enables the web-based management console.
 The no form of the command disables the web-based management console.

Syntax Description	N/A
Default	enable
Configuration Mode	Config
History	3.1.0000 3.4.0000 Updated Example
Role	admin
Example	<pre> switch (config) # web enable 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: all HTTPS certificate name: default-cert Listen enabled: yes No Listen Interfaces. 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 switch (config) # </pre>
Related Commands	show web
Note	

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 enabled HTTP TCP port is 80 HTTP redirect to HTTPS is disabled	
Configuration Mode	Config	
History	3.1.0000	
	3.4.0000	Updated Example
Role	admin	

Example

```
switch (config) # web http enable
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:     all
  HTTPS certificate name: default-cert
  Listen enabled:        yes
  No Listen Interfaces.

  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
switch (config) #
```

Related Commands

```
show web
web enable
```

Note

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

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	
	3.4.0000	Updated Example
Role	admin	
Example	<pre>switch (config) # web httpd listen enable 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: all HTTPS certificate name: default-cert Listen enabled: yes No Listen Interfaces. 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 switch (config) #</pre>	

Related Commands	N/A

Note	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.

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”
Role	admin	

Example

```
switch (config) # web https enable
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:     all
  HTTPS certificate name: default-cert
  Listen enabled:        yes
  No Listen Interfaces.

  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
switch (config) #
```

Related Commands

```
show web
web enable
```

Note

- 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
-
-

web session

web session {renewal <minutes> | timeout <minutes>}
no web session {renewal | timeout}

Configures session settings.
The no form of the command resets session settings to default.

Syntax Description	renewal <minutes>	Configures time before expiration to renew a session.
	timeout <minutes>	Configures time after which a session expires.
Default	timeout - 2.5 hours renewal - 30 min	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # web session renewal 60 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: all HTTPS certificate name: default-cert Listen enabled: yes No Listen Interfaces. Inactivity timeout: 1 hr Session timeout: 2 hr 30 min Session renewal: 60 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 switch (config) #</pre>	
Related Commands	N/A	
Note		

web proxy auth

web proxy auth {authtype <type>| basic [password <password> | username <username>]}

no web proxy auth {authtype | basic {password | username } }

Configures authentication settings for web proxy authentication.
The no form of the command resets the attributes to their default values.

Syntax Description	type	Configures the type of authentication to use with web proxy. The possible values are: • basic - HTTP basic authentication • none - No authentication
	basic	Configures HTTP basic authentication settings for proxy. The password is accepted and stored in plaintext.
	password	A password used for HTTP basic authentication with the web proxy.
	username	A username used for HTTP basic authentication with the web proxy.
Default	Web proxy is disabled.	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	

Example

```
switch (config) # web proxy auth authtype basic
switch (config) # web proxy auth basic username web-user
switch (config) # web proxy auth basic password web-password
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:    all
HTTPS certificate name: default-cert
Listen enabled:       yes
No Listen Interfaces.
```

```
Inactivity timeout: 1 hr
Session timeout:    2 hr 30 min
Session renewal:    30 min
```

Web file transfer proxy:

```
Proxy enabled: yes
Proxy address:  10.10.10.11
Proxy port:     40
Authentication type: basic
Basic auth username: web-user
Basic auth password: web-password
```

Web file transfer certificate authority:

```
HTTPS server cert verify: yes
HTTPS supplemental CA list: default-ca-list
```

```
switch (config) #
```

Related Commands

```
show web
web proxy host
```

Note

web proxy host

**web proxy host <IP address> [port <port number>]
no web proxy**

Adds and enables a proxy to be used for any HTTP or FTP downloads.
The no form of the command disables the web proxy.

Syntax Description	IP address	IPv4 or IPv6 address.
	port number	Sets the web proxy default port.
Default	1080	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # web proxy host 10.10.10.10 port 1080 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: all HTTPS certificate name: default-cert Listen enabled: yes No Listen Interfaces. Inactivity timeout: 1 hr Session timeout: 2 hr 30 min Session renewal: 30 min Web file transfer proxy: Proxy enabled: yes Proxy address: 10.10.10.10 Proxy port: 1080 Authentication type: basic Basic auth username: web-user Basic auth password: web-password Web file transfer certificate authority: HTTPS server cert verify: yes HTTPS supplemental CA list: default-ca-list switch (config) #</pre>	
Related Commands	web proxy auth	
Note		

show web

show web

Displays the web configuration.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000 3.4.0000 Updated Example 3.4.1100 Updated Example
Role	admin
Example	<pre>switch (config) # show web Web User Interface: Web interface enabled: yes Web caching enabled: yes HTTP enabled: yes HTTP port: 80 HTTP redirect to HTTPS: no HTTPS enabled: yes HTTPS port: 443 HTTPS ssl-ciphers: all HTTPS certificate name: default-cert Listen enabled: yes No Listen Interfaces. Inactivity timeout: 1 hr Session timeout: 2 hr 30 min Session renewal: 30 min Web file transfer proxy: Proxy enabled: yes Proxy address: 10.10.10.11 Proxy port: 40 Authentication type: basic Basic auth username: web-user Basic auth password: web-password Web file transfer certificate authority: HTTPS server cert verify: yes HTTPS supplemental CA list: default-ca-list switch (config) #</pre>
Related Commands	show web web proxy auth
Note	

4 System Management

4.1 Management Interface

Management interfaces are used in order to provide access to switch management user interfaces (e.g. CLI, WebUI). Mellanox switches support out-of-band (OOB) dedicated interfaces (e.g. mgmt0, mgmt1) and in-band dedicated interfaces. In addition, most Mellanox switches feature a serial port that provides access to the CLI only.

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

Step 1. Enter Config configuration mode. Run:

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

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

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

Step 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

Step 1. Enable IPv6 on this interface. Run:

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

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

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

Step 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>
Mellanox MLNX-OS Switch Management
Password:
Mellanox Switch
Mellanox 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 such 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
switch (config)#
```

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 VPI and Ethernet. It can be enabled with IP Routing but not with IP Proxy-ARP.

➤ **To set an in-band management channel:**

Step 1. Create a VLAN. Run:

Step 2. Create a VLAN interface. Run:

Step 3. Enter the VLAN interface configuration mode and configure L3 attributes. Run:

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

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 (e.g. SM HA, GW HA) the user would need to make sure the HA master has the same HA node names as the DHCP server.

4.1.7 Commands

4.1.7.1 Interface

This chapter describes the commands should be used to configure and monitor the management interface.

interface

interface {mgmt0 | mgmt1 | lo | vlan<id> | ib0}

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).
	ib0	IPoIB in-band management, relevant only for InfiniBand switch systems.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # interface mgmt0 switch (config interface mgmt0) #	
Related Commands	show interfaces <ifname>	
Notes		

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	
Role	admin	
Example	<pre> switch (config) # interface mgmt0 switch (config interface mgmt0) # ip address 10.10.10.10 255.255.255.0 switch (config interface mgmt0) # show interfaces mgmt0 Interface mgmt0 state Admin up: yes Link up: yes IP address: 10.10.10.10 Netmask: 255.255.255.0 IPv6 enabled: yes Autoconf enabled: no Autoconf route: yes Autoconf privacy: no IPv6 addresses: 1 IPv6 address: fe80:202:c9ff:fe5e:a5d8/64 Speed: 1000Mb/s (auto) Duplex: full (auto) Interface type: ethernet Interface ifindex: 2 Interface source: physical MTU: 1500 HW address: 00:02:C9:5E:A5:D8 Comment: RX bytes: 2946769856 RX packets: 44866091 RX mcast packets: 0 RX discards: 0 RX errors: 0 RX overruns: 0 RX frame: 0 TX bytes: 467577486 TX packets: 1385520 TX discards: 0 TX errors: 0 TX overruns: 0 TX carrier: 0 TX collisions: 0 TX queue len: 1000 switch (config interface mgmt0) # </pre>	

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.

ip default-gateway

ip default-gateway <next hop IP address or interface name>
no ip default-gateway

Configures a default route.
 The no form of the command removes the current default route.

Syntax Description	next hop IP address or interface name IP address, lo, mgmt0, or mgmt1.
Default	N/A
Configuration Mode	Config Interface Management
History	3.1.0000
Role	admin
Example	switch (config) # ip default-gateway mgmt1 switch (config) #
Related Commands	
Notes	

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	
Role	admin	

Example

```
switch (config interface mgmt0) # alias 2 ip address 9.9.9.9
255.255.255.255
switch (config interface mgmt0) # show interfaces mgmt0
Interface mgmt0 state
  Admin up:          yes
  Link up:           yes
  IP address:        172.30.2.2
  Netmask:           255.255.0.0
  Secondary address: 9.9.9.9/32 (alias: 'mgmt0:2')
  IPv6 enabled:      yes
  Autoconf enabled:  no
  Autoconf route:    yes
  Autoconf privacy:  no
  IPv6 addresses:    1
  IPv6 address:      fe80::202:c9ff:fe5e:a5d8/64
  Speed:             1000Mb/s (auto)
  Duplex:            full (auto)
  Interface type:    ethernet
  Interface ifindex: 2
  Interface source:  physical
  MTU:               1500
  HW address:        00:02:C9:5E:A5:D8
  Comment:

RX bytes:           2970074221      TX bytes:           468579522
RX packets:         44983023       TX packets:         1390539
RX mcast packets:   0              TX discards:        0
RX discards:        0              TX errors:           0
RX errors:          0              TX overruns:         0
RX overruns:        0              TX carrier:          0
RX frame:           0              TX collisions:       0
                                   TX queue len: 1000

switch (config interface mgmt0) #
```

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

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	
Role	admin	
Example	<pre> switch (config interface mgmt0) # mtu 1500 switch (config interface mgmt0) # show interfaces mgmt0 Interface mgmt0 state Admin up: yes Link up: yes IP address: 172.30.2.2 Netmask: 255.255.0.0 Secondary address: 9.9.9.9/32 (alias: 'mgmt0:2') IPv6 enabled: yes Autoconf enabled: no Autoconf route: yes Autoconf privacy: no IPv6 addresses: 1 IPv6 address: fe80:202:c9ff:fe5e:a5d8/64 Speed: 1000Mb/s (auto) Duplex: full (auto) Interface type: ethernet Interface ifindex: 2 Interface source: physical MTU: 1500 HW address: 00:02:C9:5E:A5:D8 Comment: </pre> RX bytes: 2970074221 TX bytes: 468579522 RX packets: 44983023 TX packets: 1390539 RX mcast packets: 0 TX discards: 0 RX discards: 0 TX errors: 0 RX errors: 0 TX overruns: 0 RX overruns: 0 TX carrier: 0 RX frame: 0 TX collisions: 0 TX queue len: 1000 <pre> switch (config interface mgmt0) # </pre>	
Related Commands	show interfaces <ifname>	
Notes		

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	
Role	admin	
Example	<pre> switch (config interface mgmt0) # duplex auto switch (config interface mgmt0) # show interfaces mgmt0 Interface mgmt0 state Admin up: yes Link up: yes IP address: 172.30.2.2 Netmask: 255.255.0.0 Secondary address: 9.9.9.9/32 (alias: 'mgmt0:2') IPv6 enabled: yes Autoconf enabled: no Autoconf route: yes Autoconf privacy: no IPv6 addresses: 1 IPv6 address: fe80::202:c9ff:fe5e:a5d8/64 Speed: 1000Mb/s (auto) Duplex: full (auto) Interface type: ethernet Interface ifindex: 2 Interface source: physical MTU: 1500 HW address: 00:02:C9:5E:A5:D8 Comment: RX bytes: 2970074221 TX bytes: 468579522 RX packets: 44983023 TX packets: 1390539 RX mcast packets: 0 TX discards: 0 RX discards: 0 TX errors: 0 RX errors: 0 TX overruns: 0 RX overruns: 0 TX carrier: 0 RX frame: 0 TX collisions: 0 TX queue len: 1000 switch (config interface mgmt0) # </pre>	

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
-
-

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		
Role	admin		
Example	<pre>switch (config interface mgmt0) # speed auto switch (config interface mgmt0) # show interfaces mgmt0 Interface mgmt0 state Admin up: yes Link up: yes IP address: 172.30.2.2 Netmask: 255.255.0.0 Secondary address: 9.9.9.9/32 (alias: 'mgmt0:2') IPv6 enabled: yes Autoconf enabled: no Autoconf route: yes Autoconf privacy: no IPv6 addresses: 1 IPv6 address: fe80::202:c9ff:fe5e:a5d8/64 Speed: 1000Mb/s (auto) Duplex: full (auto) Interface type: ethernet Interface ifindex: 2 Interface source: physical MTU: 1500 HW address: 00:02:C9:5E:A5:D8 Comment: RX bytes: 2970074221 TX bytes: 468579522 RX packets: 44983023 TX packets: 1390539 RX mcast packets: 0 TX discards: 0 RX discards: 0 TX errors: 0 RX errors: 0 TX overruns: 0 RX overruns: 0 TX carrier: 0 RX frame: 0 TX collisions: 0 TX queue len: 1000 switch (config interface mgmt0) #</pre>		

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
-

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	
Role	admin	
Example	<pre>switch (config interface mgmt0) # dhcp switch (config) # show interfaces mgmt0 configured Interface mgmt0 configuration Enabled: yes DHCP: yes Zeroconf: no IP address: Netmask: IPv6 enabled: yes Autoconf enabled: no Autoconf route: yes Autoconf privacy: no IPv6 addresses: 0 Speed: auto Duplex: auto MTU: 1500 Comment:</pre>	
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.	

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
Role	admin
Example	switch (config interface mgmt0) # dhcp hostname switch (config interface mgmt0) #
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 • After upgrading to version 3.5.1000 when a default hostname is not configured, the DHCP server assigns the new hostname for your machine • These commands do not work on in-band interfaces

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
Role	admin
Example	<pre> switch (config interface mgmt0) # no shutdown switch (config) # show interfaces mgmt0 configured Interface mgmt0 configuration Enabled: yes DHCP: yes DHCP Hostname: yes Zeroconf: no IP address: Netmask: IPv6 enabled: yes Autoconf enabled: no Autoconf route: yes Autoconf privacy: no IPv6 addresses: 0 Speed: auto Duplex: auto MTU: 1500 Comment: switch (config) # </pre>
Related Commands	show interfaces <ifname> configured
Notes	

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
Role	admin
Example	<pre>switch (config interface mgmt0) # zeroconf switch (config) # show interfaces mgmt0 configured Interface mgmt0 configuration Enabled: yes DHCP: no DHCP Hostname: yes Zeroconf: yes IP address: Netmask: IPv6 enabled: yes Autoconf enabled: no Autoconf route: yes Autoconf privacy: no IPv6 addresses: 0 Speed: auto Duplex: auto MTU: 1500 Comment:</pre>
Related Commands	show interfaces <ifname> configured
Notes	Enabling zeroconf disables DHCP and vice versa.

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		
Role	admin		
Example	<pre>switch (config interface mgmt0) # comment my-interface switch (config interface mgmt0) # show interfaces mgmt0 Interface mgmt0 state Admin up: yes Link up: yes IP address: 172.30.2.2 Netmask: 255.255.0.0 IPv6 enabled: yes Autoconf enabled: no Autoconf route: yes Autoconf privacy: no IPv6 addresses: 1 IPv6 address: fe80::202:c9ff:fe5e:a5d8/64 Speed: 1000Mb/s (auto) Duplex: full (auto) Interface type: ethernet Interface ifindex: 2 Interface source: physical MTU: 1500 HW address: 00:02:C9:5E:A5:D8 Comment: my-interface RX bytes: 962067812 TX bytes: 40658219 RX packets: 3738865 TX packets: 142345 RX mcast packets: 0 TX discards: 0 RX discards: 0 TX errors: 0 RX errors: 0 TX overruns: 0 RX overruns: 0 TX carrier: 0 RX frame: 0 TX collisions: 0 TX queue len: 1000 switch (config interface mgmt0) #</pre>		

Related Commands	N/A

Notes

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
Role	admin
Example	<pre> switch (config interface mgmt0) # ipv6 enable switch (config interface mgmt0) # show interfaces mgmt0 Interface mgmt0 state Admin up: yes Link up: yes IP address: 172.30.2.2 Netmask: 255.255.0.0 IPv6 enabled: yes Autoconf enabled: no Autoconf route: yes Autoconf privacy: no IPv6 addresses: 1 IPv6 address: fe80::202:c9ff:fe5e:a5d8/64 Speed: 1000Mb/s (auto) Duplex: full (auto) Interface type: ethernet Interface ifindex: 2 Interface source: physical MTU: 1500 HW address: 00:02:C9:5E:A5:D8 Comment: my-interface RX bytes: 962067812 TX bytes: 40658219 RX packets: 3738865 TX packets: 142345 RX mcast packets: 0 TX discards: 0 RX discards: 0 TX errors: 0 RX errors: 0 TX overruns: 0 RX overruns: 0 TX carrier: 0 RX frame: 0 TX collisions: 0 TX queue len: 1000 switch (config interface mgmt0) # </pre>

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.
-
-

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	
Role	admin	

Example

```
switch (config interface mgmt0) # ipv6 fe80::202:c9ff:fe5e:a5d8/64
switch (config interface mgmt0) # show interfaces mgmt0
Interface mgmt0 state
  Admin up:          yes
  Link up:           yes
  IP address:        172.30.2.2
  Netmask:           255.255.0.0
  IPv6 enabled:      yes
  Autoconf enabled:  no
  Autoconf route:    yes
  Autoconf privacy:  no
  IPv6 addresses:    1
  IPv6 address:      fe80::202:c9ff:fe5e:a5d8/64
  Speed:             1000Mb/s (auto)
  Duplex:            full (auto)
  Interface type:    ethernet
  Interface ifindex: 2
  Interface source:  physical
  MTU:               1500
  HW address:        00:02:C9:5E:A5:D8
  Comment:           my-interface

  RX bytes:          962067812      TX bytes:          40658219
  RX packets:        3738865       TX packets:        142345
  RX mcast packets:  0             TX discards:       0
  RX discards:       0             TX errors:         0
  RX errors:         0             TX overruns:       0
  RX overruns:       0             TX carrier:        0
  RX frame:          0             TX collisions:     0
                                   TX queue len:    1000

switch (config interface mgmt0) #
```

Related Commands

```
ipv6 enable
show interface <ifname>
```

Notes

- Unlike IPv4, IPv6 can have multiple IPv6 addresses on a given interface
- For Ethernet, the default interface identifier is a 64-bit long modified EUI-64, which is based on the MAC address of the interface

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	
Role	admin	
Example	switch (config) # ipv6 dhcp primary-intf mgmt0 switch (config) #	
Related Commands	ipv6 enable ipv6 address show interface <ifname>	
Notes		

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
Role	admin
Example	switch (config) # ipv6 dhcp stateless switch (config) #
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

show interface brief

show interface <ifname> brief

Displays a brief info on the interface configuration and status.

Syntax Description	ifname	The interface name e.g., “mgmt0”, “mgmt1”, “lo” (loopback), etc.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # show interfaces mgmt0 brief Interface mgmt0 state Admin up: yes Link up: yes IP address: 172.30.2.2 Netmask: 255.255.0.0 IPv6 enabled: yes Autoconf enabled: no Autoconf route: yes Autoconf privacy: no IPv6 addresses: 1 IPv6 address: fe80::202:c9ff:fe5e:a5d8/64 Speed: 1000Mb/s (auto) Duplex: full (auto) Interface type: ethernet Interface ifindex: 2 Interface source: physical MTU: 1500 HW address: 00:02:C9:5E:A5:D8 Comment: my-interface switch (config) #</pre>	
Related Commands	N/A	
Notes		

show interface configured

show interface <ifname> configured

Displays configuration information about the specified interface.

Syntax Description	ifname	The interface name e.g., “mgmt0”, “mgmt1”, “lo” (loopback), etc.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.1.0000	
	3.5.1000	Updated Example with “DHCP Hostname”
Role	admin	
Example	<pre>switch (config) # show interfaces mgmt0 configured Interface mgmt0 configuration Enabled: yes DHCP: yes DHCP Hostname: yes Zeroconf: no IP address: Netmask: IPv6 enabled: yes Autoconf enabled: no Autoconf route: yes Autoconf privacy: no IPv6 addresses: 0 Speed: auto Duplex: auto MTU: 1500 Comment: my-interface</pre>	
Related Commands	N/A	
Notes		

4.1.7.2 Hostname Resolution

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 “.”
Role	admin	
Example	switch (config) # hostname my-switch-hostname my-switch-hostname (config) #	
Related Commands	show hosts	
Notes	• Hostname may contain letters, numbers, and hyphens ('-'), in any combination • Hostname may not contain other characters, such as “%”, “_” etc. • The character “.” is supported • Hostname may not begin with a hyphen • Hostname may be 1-63 characters long • Changing hostname stamps a new HTTPS certificate	

ip name-server

ip name-server <IPv4/IPv6 address>
no name-server <IPv4/IPv6 address>

Sets the static name server.
The no form of the command clears the name server.

Syntax Description	IPv4/v6 address	IPv4 or IPv6 address.
Default	No server name	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ip name-server 9.9.9.9 switch (config) # show hosts Hostname: switch Name server: 9.9.9.9 (configured) Name server: 10.211.0.121 (dynamic) Name server: 172.30.0.126 (dynamic) Name server: 10.4.0.135 (dynamic) Domain name: lab.mtl.com (dynamic) Domain name: vmlab.mtl.com (dynamic) Domain name: yok.mtl.com (dynamic) Domain name: mtl.com (dynamic) IP 127.0.0.1 maps to hostname localhost IPv6 ::1 maps to hostname localhost6 Automatically map hostname to loopback address: yes Automatically map hostname to IPv6 loopback address: no switch (config) #</pre>	
Related Commands	show hosts	
Notes		

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. Domain names are formed by the rules and procedures of the Domain Name System (DNS).
Default	No static domain name	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ip domain-list mydomain.com switch (config) # show hosts Hostname: switch Name server: 10.211.0.121 (dynamic) Name server: 172.30.0.126 (dynamic) Name server: 10.4.0.135 (dynamic) Domain name: mydomain.com (configured) Domain name: lab.mtl.com (dynamic) Domain name: vmlab.mtl.com (dynamic) Domain name: yok.mtl.com (dynamic) Domain name: mtl.com (dynamic) IP 1.1.1.1 maps to hostname p IP 127.0.0.1 maps to hostname localhost IPv6 ::1 maps to hostname localhost6 Automatically map hostname to loopback address: yes Automatically map hostname to IPv6 loopback address: no switch (config) #</pre>	
Related Commands	show hosts	
Notes		

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	
Role	admin	
Example	<pre>switch (config) # ip host my-host 2.2.2.2 switch (config) # ipv6 host my-ipv6-host 2001::8f9 switch (config) # show hosts Hostname: switch Name server: 9.9.9.9 (configured) Name server: 10.211.0.121 (dynamic) Name server: 172.30.0.126 (dynamic) Name server: 10.4.0.135 (dynamic) Domain name: mydomain.com (configured) Domain name: lab.mtl.com (dynamic) Domain name: vmlab.mtl.com (dynamic) Domain name: yok.mtl.com (dynamic) Domain name: mtl.com (dynamic) IP 1.1.1.1 maps to hostname p IP 127.0.0.1 maps to hostname localhost IP 2.2.2.2 maps to hostname my-host IPv6 2001::8f9 maps to hostname my-ipv6-host IPv6 ::1 maps to hostname localhost6 Automatically map hostname to loopback address: yes Automatically map hostname to IPv6 loopback address: yes switch (config) #</pre>	
Related Commands	show hosts	
Notes		

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
Role	admin
Example	<pre> switch (config) # ip map-hostname switch (config) # # show hosts Hostname: switch Name server: 9.9.9.9 (configured) Name server: 10.211.0.121 (dynamic) Name server: 172.30.0.126 (dynamic) Name server: 10.4.0.135 (dynamic) Domain name: mydomain.com (configured) Domain name: lab.mtl.com (dynamic) Domain name: vmlab.mtl.com (dynamic) Domain name: yok.mtl.com (dynamic) Domain name: mtl.com (dynamic) IP 1.1.1.1 maps to hostname p IP 127.0.0.1 maps to hostname localhost IP 2.2.2.2 maps to hostname my-host IPv6 2001::8f9 maps to hostname my-ipv6-host IPv6 ::1 maps to hostname localhost6 Automatically map hostname to loopback address: yes Automatically map hostname to IPv6 loopback address: yes switch (config) # switch (config) # ping my-host-name PING localhost (127.0.0.1) 56(84) bytes of data. 64 bytes from localhost (127.0.0.1): icmp_seq=1 ttl=64 time=0.078 ms 64 bytes from localhost (127.0.0.1): icmp_seq=2 ttl=64 time=0.052 ms 64 bytes from localhost (127.0.0.1): icmp_seq=3 ttl=64 time=0.058 ms </pre>

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.
-
-

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
Role	admin
Example	<pre>switch (config) # show hosts Hostname: my-host-name Name server: 9.9.9.9 (configured) Name server: 10.211.0.121 (dynamic) Name server: 172.30.0.126 (dynamic) Name server: 10.4.0.135 (dynamic) Domain name: mydomain.com (configured) Domain name: lab.mtl.com (dynamic) Domain name: vmlab.mtl.com (dynamic) Domain name: yok.mtl.com (dynamic) Domain name: mtl.com (dynamic) IP 1.1.1.1 maps to hostname p IP 127.0.0.1 maps to hostname localhost IP 2.2.2.2 maps to hostname my-host IPv6 ::1 maps to hostname localhost6 Automatically map hostname to loopback address: yes Automatically map hostname to IPv6 loopback address: no switch (config) #</pre>
Related Commands	N/A
Notes	

4.1.7.3 Routing

ip/ipv6 route

{ip | ipv6} route vrf <vrf-name> <network-prefix> <netmask> <next-hop>
no ip route <vrf-name> <network-prefix> <netmask> <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).																																	
Default	N/A																																		
Configuration Mode	Config																																		
History	3.1.0000																																		
Role	admin																																		
Example	<pre>switch (config) # ip route 20.20.20.0 255.255.255.0 mgmt0 switch (config) # show ip route</pre> <table><tr><th>Destination</th><th>Mask</th><th>Gateway</th><th>Interface</th><th>Source</th></tr><tr><td>default</td><td>0.0.0.0</td><td>172.30.0.1</td><td>mgmt0</td><td>DHCP</td></tr><tr><td>10.10.10.10</td><td>255.255.255.255</td><td>0.0.0.0</td><td>mgmt0</td><td>static</td></tr><tr><td>20.10.10.10</td><td>255.255.255.255</td><td>172.30.0.1</td><td>mgmt0</td><td>static</td></tr><tr><td>20.20.20.0</td><td>255.255.255.0</td><td>0.0.0.0</td><td>mgmt0</td><td>static</td></tr><tr><td>172.30.0.0</td><td>255.255.0.0</td><td>0.0.0.0</td><td>mgmt0</td><td>interface</td></tr></table>					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
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																															
Related Commands	show ip route																																		
Notes																																			

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	First version
	3.2.0500	removed IPv4 configuration option
Role	admin	
Example	<pre>switch (config) # ipv6 default-gateway ::1 switch (config) # show ipv6 default-gateway static Configured default gateways: ::1 switch (config) #</pre>	
Related Commands	show ip route	
Notes	- The configured default gateway will not be used if DHCP is enabled. - In order to configure ipv4 default-gateway use 'ip route' command.	

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	
Role	admin	
Example	<pre> 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 switch (config) # </pre>	
Related Commands	show ip default-gateway	
Notes		

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	
Role	admin	
Example	<pre>switch (config) # ipv6 default-gateway 10.10.10.10 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</pre>	
Related Commands	ipv6 default-gateway	
Notes	The configured IPv4 default gateway will not be used if DHCP is enabled.	

4.1.7.4 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.

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																																			
Role	admin																																			
Example	<pre>switch (config interface mgmt0) #ip arp 20.20.20.20 aa:aa:aa:aa:aa:aa switch (config interface mgmt0) # show ip arp</pre> Total number of entries: 6 <table><thead><tr><th>Address</th><th>Type</th><th>MAC Address</th><th>Interface</th></tr></thead><tbody><tr><td>10.209.1.103</td><td>Dynamic</td><td>00:02:C9:11:A1:78</td><td>mgmt0</td></tr><tr><td>10.209.1.168</td><td>Dynamic</td><td>00:02:C9:5E:C3:28</td><td>mgmt0</td></tr><tr><td>10.209.1.104</td><td>Dynamic</td><td>00:02:C9:11:A1:E6</td><td>mgmt0</td></tr><tr><td>10.209.1.153</td><td>Dynamic</td><td>00:02:C9:11:A1:86</td><td>mgmt0</td></tr><tr><td>10.209.1.105</td><td>Dynamic</td><td>00:02:C9:5E:0B:56</td><td>mgmt0</td></tr><tr><td>10.209.0.1</td><td>Dynamic</td><td>00:00:5E:00:01:01</td><td>mgmt0</td></tr><tr><td>20.20.20.20</td><td>Static</td><td>AA:AA:AA:AA:AA:AA</td><td>mgmt0</td></tr></tbody></table> <pre>switch (config interface mgmt0) #</pre>				Address	Type	MAC Address	Interface	10.209.1.103	Dynamic	00:02:C9:11:A1:78	mgmt0	10.209.1.168	Dynamic	00:02:C9:5E:C3:28	mgmt0	10.209.1.104	Dynamic	00:02:C9:11:A1:E6	mgmt0	10.209.1.153	Dynamic	00:02:C9:11:A1:86	mgmt0	10.209.1.105	Dynamic	00:02:C9:5E:0B:56	mgmt0	10.209.0.1	Dynamic	00:00:5E:00:01:01	mgmt0	20.20.20.20	Static	AA:AA:AA:AA:AA:AA	mgmt0
Address	Type	MAC Address	Interface																																	
10.209.1.103	Dynamic	00:02:C9:11:A1:78	mgmt0																																	
10.209.1.168	Dynamic	00:02:C9:5E:C3:28	mgmt0																																	
10.209.1.104	Dynamic	00:02:C9:11:A1:E6	mgmt0																																	
10.209.1.153	Dynamic	00:02:C9:11:A1:86	mgmt0																																	
10.209.1.105	Dynamic	00:02:C9:5E:0B:56	mgmt0																																	
10.209.0.1	Dynamic	00:00:5E:00:01:01	mgmt0																																	
20.20.20.20	Static	AA:AA:AA:AA:AA:AA	mgmt0																																	
Related Commands	<pre>show ip arp ip route</pre>																																			
Notes																																				

ip arp timeout

ip arp [vrf <vrf-name>] timeout <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
Role	admin	
Example	switch (config) # ip arp timeout 2000 switch (config) #	
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	

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	
Role	admin	
Example	<pre>switch-626a54 [standalone: master] (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) switch (config) #</pre>	
Related Commands		
Notes		

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	
Role	admin	
Example	<pre>switch (config) # ipv6 neighbor 2001:db8:701f::8f9 mgmt0 00:11:22:33:44:55 switch (config) #</pre>	
Related Commands	<pre>show ipv6 neighbor ipv6 route arp clear ipv6 neighbors</pre>	
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.	

clear ipv6 neighbors

clear ipv6 neighbors

Clears the dynamic neighbors cache.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # clear ipv6 neighbors switch (config) #
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

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	Config																				
History	3.1.0000																				
Role	admin																				
Example	switch (config) # show ipv6 neighbors <table><thead><tr><th>IPv6 Address</th><th>Age</th><th>MAC Address</th><th>State</th><th>Interf</th></tr></thead><tbody><tr><td colspan="5">-----</td></tr><tr><td>2001::2</td><td>9428</td><td>AA:AA:AA:AA:AA:AA</td><td>permanent</td><td>mgmt0</td></tr></tbody></table> switch (config) #						IPv6 Address	Age	MAC Address	State	Interf	-----					2001::2	9428	AA:AA:AA:AA:AA:AA	permanent	mgmt0
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																					

4.1.7.5 DHCP

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-host-
name}
```

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.
	send-hostname	Enables the DHCP client to send a hostname during negotiation.
Default	no ip dhcp yield-to-static no ip dhcp hostname ip ip dhcp primary-intf mgmt0 no ip dhcp send-hostname	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	

Example

```
switch (config) # ip dhcp default-gateway yield-to-static
switch (config) # show ip dhcp
```

Interface	DHCP Enabled	DHCP Running	Valid lease
lo	no	no	no
mgmt0	yes	yes	yes
mgmt1	yes	yes	no

```
DHCP primary interface:
  Configured: mgmt0
  Active:     mgmt0

DHCP default gateway yields to static configuration: yes

DHCP client options:
  Send Hostname: no
  Client Hostname: switch (using system hostname)
switch (config) #
```

Related Commands

```
show ip dhcp
dhcp [renew]
```

Notes

DHCP is supported for IPv4 networks only.

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
Role	admin
Example	<pre>switch (config) # show ip dhcp DHCP primary interface: Configured: mgmt0 Active: mgmt0 DHCP: yield default gateway to static configuration: yes DHCP Client Options: Send Hostname: no Client Hostname: switch (using system hostname) switch (config) #</pre>
Related Commands	<pre>ip dhcp dhcp [renew]</pre>
Notes	

4.1.7.6 General IPv6 Commands

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
Role	admin
Example	<pre>switch (config) # ipv6 enable switch (config) # show ipv6 IPv6 summary IPv6 supported: yes IPv6 admin enabled: yes IPv6 interface count: 2 switch (config) #</pre>
Related Commands	ipv6 default-gateway ipv6 host ipv6 map-hostname ipv6 neighbor ipv6 route show ipv6 show ipv6 default-gateway show ipv6 route
Notes	

4.1.7.7 IP Diagnostic Tools

ping

ping [-LRUbdnqrVvA] [-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 http://linux.about.com/od/commands/l/blcm-dl8_ping.htm
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>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 switch (config) #</pre>
Related Commands	tracert
Notes	

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 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 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.
	-q	Sets the number of probes per each hop. Default is 3.

Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # traceroute 192.168.10.70 traceroute to 192.168.10.70 (192.168.10.70), 30 hops max, 40 byte pack- ets 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 switch (config) #</pre>
Related Commands	
Notes	

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
Role	admin
Example	<pre>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 switch (config) #</pre>
Related Commands	N/A
Notes	

4.2 NTP, Clock & Time Zones

Network Time Protocol (NTP) is a networking protocol for clock synchronization between computer systems over packet-switched, 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.

For an example, please refer to “[HowTo enable NTP on Mellanox switches](https://community.mellanox.com)” in the Mellanox Community (<https://community.mellanox.com>).

4.2.1 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***

Step 1. Create the key. Run:

```
switch (config)# ntp authentication-key 1 md5 password
```

Step 2. Add the key to the trusted list. Run:

```
switch (config)# ntp trusted-key 1
```

Step 3. Assign the key to the server/peer. Run:

```
switch (config)# ntp server 10.34.1.1 keyID 1
```

4.2.2 NTP Authentication Key

An authentication key may be created and used to authenticate incoming NTP packets.

For the key to be used:

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 switch.

4.2.3 Commands

clock set

clock set <hh:mm:ss> [<yyyy/mm/dd>]

Sets the time and date.

Syntax Description	hh:mm:ss	Time.
	yyyy/mm/dd	Date.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # clock set 23:23:23 2010/08/19 switch (config) # show clock Time: 23:23:26 Date: 2010/08/19 Time zone: UTC (Etc/UTC) UTC offset: same as UTC switch (config) #</pre>	
Related Commands	show clock	
Notes	If not specified, the date will be left the same.	

clock timezone

clock timezone [<zone word> [<zone word> [<zone word>] [<zone word>]]

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	The 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	
Role	admin	
Example	<pre>switch (config) # clock timezone America North United_States Other New_York switch (config) # show clock Time: 10:08:53 Date: 2015/10/29 Time zone: America North United_States Other New_York (America/New_York) UTC offset: -0400 (UTC minus 4 hours) switch (config) #</pre>	
Related Commands	show clock	
Notes		

ntp

ntp {disable | enable | {peer | server} <IP address> [version <number> | disable]}
no ntp {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 or 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
Default	NTP is enabled NTP version number is 4	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # no ntp peer 192.168.10.24 disable switch (config) #	
Related Commands	N/A	
Notes		

ntpdate

ntpdate <IP address>

Sets the system clock using the specified SNTP server.

Syntax Description	IP address	IP.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>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 switch (config) #</pre>	
Related Commands	N/A	
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.	

ntp authenticate

ntp authenticate
no ntp authenticate

Enables NTP authentication.
The no form of the command disables NTP authentication.

Syntax Description	N/A	N/A
Default	Disabled	
Configuration Mode	Config	
History	3.5.0200	
Role	admin	
Example	switch (config) # ntp authenticate	
Related Commands	N/A	
Notes		

ntp authentication-key

ntp authentication-key <key_id> <encrypt_type> [<password>]
no ntp authentication-key <key_id>

Adds a new authentication key and stores it.
The no form of the command removes key ID configuration if it exists.

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	
Role	admin	
Example	<pre>switch (config) # ntp authentication-key 123 md5 examplepass switch (config) # ntp authentication-key 1234 sha1 Password: ** Confirm: ** switch (config) #</pre>	
Related Commands	N/A	
Notes	If a password is not entered, a prompt appears requiring that a password is introduced.	

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 and IPv6 are acceptable)
Default	Disabled	
Configuration Mode	Config	
History	3.5.0200	
Role	admin	
Example	switch (config) # ntp peer 10.10.10.10 disable switch (config) #	
Related Commands	N/A	
Notes		

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 and IPv6 are acceptable)
	key_id	Range: 1-65534
Default	Disabled	
Configuration Mode	Config	
History	3.5.0200	
Role	admin	
Example	switch (config) # ntp peer 10.10.10.10 keyID 120	
Related Commands	N/A	
Notes		

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 and IPv6 are acceptable)
	ver_num	NTP version (3 or 4)
Default	4	
Configuration Mode	Config	
History	3.5.0200	
Role	admin	
Example	switch (config) # ntp peer 10.10.10.10 version 4	
Related Commands	N/A	
Notes		

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 server (IPv4 and IPv6 are acceptable)
Default	Disabled	
Configuration Mode	Config	
History	3.5.0000	
Role	admin	
Example	switch (config) # ntp server 10.10.10.10 disable switch (config) #	
Related Commands	N/A	
Notes		

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 server (IPv4 and IPv6 are acceptable)
	key_id	Range: 1-65534
Default	Disabled	
Configuration Mode	Config	
History	3.5.0200	
Role	admin	
Example	switch (config) # ntp server 10.10.10.10 keyID 120	
Related Commands	N/A	
Notes		

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 NTP server
Default	N/A	
Configuration Mode	Config	
History	3.6.2002	
Role	admin	
Example	switch (config) # ntp server 10.10.10.10 trusted-enable	
Related Commands	N/A	
Notes	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.	

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 server (IPv4 and IPv6 are acceptable)
	ver_num	NTP version (3 or 4)
Default	4	
Configuration Mode	Config	
History	3.5.0200	
Role	admin	
Example	switch (config) # ntp server 10.10.10.10 version 4	
Related Commands	N/A	
Notes		

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	
Role	admin	
Example	switch (config) # ntp trusted-key 1,3,5 switch (config) # ntp trusted-key 1-5	
Related Commands	ntp authentication-key	
Notes	Keys may be separated with commas without any space, or they may be set as a range using a hyphen.	

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
Role	admin
Example	<pre>switch (config) # show clock Time: 04:21:44' Date: 2012/02/26 Time zone: America North United_States Other New_York switch (config) #</pre>
Related Commands	N/A
Notes	

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
Role	admin
Example	<pre> switch (config) # show ntp NTP is administratively enabled. NTP Authentication is administratively disabled. Clock is synchronized. Reference: 10.134.46.4. Offset: -9.605 ms. Active servers and peers: 10.1.1.1 Conf Type : dual Status : pending Stratum : 16 Offset(msec) : 0.000 Ref clock : .INIT. Poll Interval (sec): 64 Last Response (sec): N/A Auth state : none 10.134.46.4 Conf Type : serv Status : sys.peer(*) Stratum : 4 Offset(msec) : -9.605 Ref clock : 10.7.77.134 Poll Interval (sec): 64 Last Response (sec): 55 Auth state : none switch (config) # </pre>
Related Commands	N/A
Notes	

show ntp configured

show ntp configured

Displays NTP configuration.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.5.0200
Role	admin
Example	<pre>switch (config) # show ntp configured NTP enabled: yes NTP Authentication enabled: no No NTP peers configured. NTP server 10.10.10.10 Enabled: yes NTP version: 4 Key ID: none switch (config) #</pre>
Related Commands	N/A
Notes	

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
Role	admin
Example	<pre>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 switch (config) #</pre>
Related Commands	N/A
Notes	

4.3 Unbreakable Links

MLNX-OS® offers a phy profile configuration for VPI interfaces.

PHY profile includes Link Level Retransmission (LLR) configuration. A PHY profile is bind to any VPI interface.

4.3.1 Link Level Retransmission (LLR)

Link Level Retransmission (LLR) is used on signal integrity marginal systems to decrease and/or eliminate the impact of physical errors on the system's performance.

- LLR transmitter breaks the transmitted Layer 2 data stream into Cells and adds a CRC checksum to each cell.
- LLR receiver checks the Cell CRC, in case there is no CRC errors, it forwards the cell and acknowledges the peer.

If a cell is dropped by the receiver the transmitter retransmits the cell.



LLR is a Mellanox proprietary feature and will only work with Mellanox to Mellanox ports.



LLR is not operational for cables longer then 30m.

LLR Mode

The following LLR modes are applicable per port per speed:

- disable – no LLR
- enable – the port becomes passive, only if it got a request to use LLR it activates, otherwise it remains disabled
- enable-request – the port becomes active, it keeps sending LLR requests to the peer

LLR Negotiation

Both ports on the link perform LLR discovery and negotiation. In order the LLR to be in active state on the link, the following should apply:

- One port must be configured with LLR “enable-request” on the specified speed.
- The other port (peer) may be configured with LLR “enable-request” or “enable” on the same specified speed



If both the local port and remote port configured with LLR “enabled” the LLR negotiation will not be activated - the ports will remain in LLR in-active state.

LLR Status

LLR status is a port parameter that states the current state of the LLR.

- Active – LLR is operationally running
- In-Active – LLR is not running

4.3.2 Configuring Phy Profile & LLR

➤ *To configure a phy profile:*

Step 1. Create/edit a phy profile and enter a phy profile configuration mode. Run:

```
switch (config) # phy-profile my-profile
switch (config phy profile my-profile) #
```

Step 2. Configure LLR attributes. Run:



All ports mapped to the phy profile must be in shutdown state before editing the profile.

```
switch (config phy profile my-profile) # llr support ib speed FDR enable-request
switch (config phy profile my-profile) # llr support ib speed QDR disable
switch (config phy profile my-profile) # ...
```

Step 3. Bind the profile to the desired interface. Run:



The port must be in shutdown state before binding the phy-profile.

```
switch (config) # interface ib 1/1
switch (config interface ib 1/1) # shutdown
switch (config interface ib 1/1) # phy-profile map my-profile
switch (config interface ib 1/1) # no shutdown
switch (config interface ib 1/1) #
```

Step 4. Verify LLR configuration and status. Run:

```
switch (config) # show interfaces ib llr

Interface phy-profile LLR status
...
ib 1/1    my-profile  Active
ib 1/2    disable     Inactive
...
switch (config) #
```

Step 5. Display phy-profile configuration. Run:

```
switch (config) # show phy-profile my-profile
Profile: my-profile
llr support ib-speed
SDR: disable
DDR: disable
QDR: disable
FDR10: enable-request
FDR: enable-request
switch (config) #
```

4.3.3 Commands

phy-profile

phy-profile <profile-name>
no phy-profile <profile-name>

Creates a PHY profile (port physical parameters), and enter the profile configuration mode.

The no form of the command deletes the phy-profile

Syntax Description	profile-name	40-byte-string.
Default	“high-speed-ber”: FDR and FDR10 speeds are LLR enable-request state, all the rest speed options are in disable state.	
Configuration Mode	Config	
History	3.2.0700	First version
	3.3.3000	Default updated
Role	admin	
Example	switch (config) # phy-profile my-profile switch (config phy-profile my-profile) #	
Related Commands		
Notes	10 profiles is the maximum profiles supported. - When deleting a profile, all interface related to that profile need to be in shutdown state.	

llr support ib-speed

llr support ib-speed <speed-options> <speed-actions>
no llr support ib-speed <speed-options>

Sets LLR InfiniBand supported speeds.
The no form of the command disables the llr on this speed.

Syntax Description	speed-options	• sdr • ddr • qdr • fdr10 • fdr
	speed-action	enable: only enable bit is on (passive mode) enable-request: both enable and request bits are on (active mode)
Default	N/A	
Configuration Mode	Config Phy-Profile	
History	3.2.0700	
Role	admin	
Example	<pre>switch (config) # phy-profile my-profile switch (config phy-profile my-profile) # llr support speed fdr enable switch (config phy-profile my-profile) #</pre>	
Related Commands		
Notes		

interface ib internal phy-profile enable llr64

interface ib internal phy-profile enable llr64
no interface ib internal phy-profile enable llr64

Enables LLR64 on the internal interfaces of director switch systems.
 The no form of the command disables LLR64 on the internal interfaces of director switch systems.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config
History	3.4.1854
Role	admin
Example	switch (config) # interface ib internal phy-profile enable llr64 Please save configuration and reboot the system for the changes to take effect. switch (config) #
Related Commands	show interfaces ib internal leaf capabilities
Notes	Running the command “show interfaces ib internal leaf capabilities” shows whether LLR64 is configured

phy-profile map

phy-profile map <profile-name>
no phy-profile map

Binds a phy-profile to the interface.
The no form of the command set the port mapping to the default profile.

Syntax Description	profile-name	40-byte-string.
Default	Default profile - “high-speed-ber” with the following attributes: SDR: disable DDR: disable QDR: disable FDR10: enable-request FDR: enable-request	
Configuration Mode	Config Interface IB	
History	3.2.0700	First version
	3.3.3000	Default updated
Role	admin	
Example	switch (config) # interface ib 1/1 switch (config interface ib 1/1) #phy-profile map my-profile switch (config interface ib 1/1) #	
Related Commands		
Notes		

show phy-profile

show phy-profile [profile-name]

Shows phy-profile list

Syntax Description	profile-name	40-byte-string. Shows a specific profile.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.2.0700	First version
	3.3.3000	Output updated.
Role	admin	
Example	<pre>switch (config) # show phy-profile Profile: high-speed-ber ----- llr support ib-speed SDR: disable DDR: disable QDR: disable FDR10: enable-request FDR: enable-request switch (config) #</pre>	
Related Commands	phy-profile	
Notes		

show interfaces ib llr

show interfaces ib [<number>] llr

Displays LLR status

Syntax Description	number	The interface number
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.2.0500	
Role	admin	
Example	<pre>switch (config) # show interfaces ib llr Interface phy-profile LLR status Ib 1/1 high-speed-ber Active Ib 1/2 high-speed-ber Inactive Ib 1/3 high-speed-ber Inactive ... switch (config) #</pre>	
Related Commands		
Notes		

4.4 Virtual Protocol Interconnect (VPI)

Virtual Protocol Interconnect (VPI) technology allows InfiniBand and Ethernet traffic to co-exist on one platform.

VPI technology can be achieved in different levels:

- System level VPI – it can be decided, per system, whether to use InfiniBand or Ethernet for all the interfaces in the system. Either Ethernet switch or InfiniBand switch profile can be configured on the system in order to determine the running link protocol for all the system ports.
- Interface level VPI – it can be decided, per system port, whether to use InfiniBand or Ethernet as a link protocol. A single VPI SWID can be configured as the system profile, and, per port, the link protocol may be selected. Configuring the switch to VPI mode allows splitting the hardware into two separate switches (an Ethernet switch and an InfiniBand switch). Traffic does not pass between those switches. While configuring the VPI system profile, bridging (or gateway) capabilities can be added to pass traffic from the Ethernet to the InfiniBand hosts.

Configuring your system to VPI single-switch mode splits your network interfaces to two groups:

- The Ethernet set of ports, which are connected to the Ethernet switch
- The InfiniBand set of ports, which are connected to the InfiniBand switch



VPI single switch profile is not a gateway. Ethernet traffic does not pass to the InfiniBand ports and vice versa.



VPI mode requires using either a SX6036G system, or installing a license (UPGR-XXXX-GW) on SX1012, SX1700, SX1710, SX1036, SX6012, SX6018, and SX6710 and SX6036. Refer to [Section 2.4, “Licenses,”](#) on page 63 for more details on the licenses.

In order to set your system to work with VPI, the system profile should be changed to “vpi-single-switch”. In addition, the required set of ports should be changed from InfiniBand to Ethernet or vice versa.

The following systems can be configured as VPI switches:

- SX1012, SX1700, SX1710, SX1036
- SX6012, SX6018, SX6710, SX6036, SX6036G



The SX6036G system supports VPI by default, with the port configured as follows:

- Interfaces 1/1-1/8 Ethernet
- Interfaces 1/9-1/36 InfiniBand

➤ **To make an Ethernet switch system support VPI in a single-switch mode:**

Step 1. Make sure you have the latest software version installed.

Step 2. Install a gateway license.

Step 3. Set the system profile to be “vpi-single-switch”.



When changing to a VPI profile the number of unicast MAC addresses is decreased by 2k to 46k entries.

Step 4. Use the `port type force` command to change the disabled ports from Ethernet to InfiniBand.



This step may take several minutes.

```
switch (config)# license install <license>
switch (config)# system profile vpi-single-switch
...
switch (config)# port 1/9-1/36 type infiniband force
switch (config)# show ports type
Ethernet: 1/1, 1/2, ... 1/8
Infiniband: 1/9, 1/10 ... 1/36
switch (config) #
```

➤ **To make an InfiniBand switch system support VPI in a single-switch mode:**

Step 1. Make sure you have the latest software version installed

Step 2. Install a gateway license. See [Section 2.4, “Licenses,” on page 63](#).

Step 3. Set the system profile to be “vpi-single-switch”.

Step 4. Use the command `port type force` to change the disabled ports from InfiniBand to Ethernet.

```
switch (config)# license install <license>
switch (config)# system profile vpi-single-switch
...
switch (config)# port 1/1-1/8 type ethernet force
switch (config)# interface ethernet 1/1-1/8 no shutdown
switch (config)# show ports type
Ethernet: 1/1, 1/2, ... 1/8
Infiniband: 1/9, 1/10 ... 1/36
switch (config) #
```



Changing the system profile deletes all the existing switch configurations and reboots the system. Management connectivity, however, is kept.



Note that the `port type force` command is valid for admin state only. In case of ETH ports, the user must remove specific configurations such as LAG, port mirror and split before moving to InfiniBand mode.

4.4.1 Commands

port type

port <slot>/<port>[-<slot>/<port>] type <ethernet/infiniband> [force]
no port <slot>/<port>[-<slot>/<port>] type <ethernet/infiniband>

Sets the port link protocol type on a specific port or a range of ports.
 The no form of the command sets the port type to default on the specified port(s).

Syntax Description	slot/port	The port number.
	type <ethernet/infiniband>	The desired port type. Options are: • InfiniBand • Ethernet
	force	Forces a port type change regardless of the admin's state
Default	The default value depends on the system: • SX10xx systems have Ethernet as default • SX60xx systems have InfiniBand as default • The SX6036G system has 1/1-1/8 as Ethernet and 1/9-1/36 as InfiniBand by default	
Configuration Mode	Config	
History	3.2.1100	
	3.3.4100	Ability to configure specific ports
	3.3.5006	Removed “force” parameter
Role	admin	
Example	<pre>switch (config) # port 1/1-1/8 type ethernet switch (config) # show ports type Ethernet: 1/1, 1/2, ... 1/8 Infiniband: 1/9, 1/10 ... 1/36 switch (config) #</pre>	

Related Commands**Notes**

- System profile must be vpi-single-switch. Refer to the command ‘system profile’.
 - For the “non-force” version of the command: The interface(s) must be disabled, and must not be split, mirrored or part of a port-channel when running the command.
 - For the “force” version of the command: The interface must not be split, mirrored or part of a port-channel when running the command.
 - If one or more ports in a range cannot change protocol, a printout indicates it but action continues on the other ports
-
-

show ports type

show ports type

Displays the link protocol configuration in the system.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.2.1100
Role	admin
Example	switch (config) # show ports type Ethernet: 1/1, 1/2, ... 1/24 Infiniband: 1/25, 1/26 ... 1/36 switch (config) #
Related Commands	
Notes	

4.5 System Profile

MLNX-OS has the ability to change the system profile upon license acquisition. The following are the possible system profiles:

- Ethernet single switch
- InfiniBand single switch
- InfiniBand single switch without adaptive routing
- VPI single switch

Changing the system profile requires one of the following licenses:

- Ethernet systems require a UPGR-10xx-GW license to change the system's profile to InfiniBand or VPI
- InfiniBand systems require a UPGR-60xx-GW license to change the system's profile to Ethernet or VPI



Changing the system profile will delete all the existing switch configuration and reboot the system. Management connectivity, however, will be kept.



Externally managed InfiniBand switch systems cannot run an Ethernet license.

➤ *To change the system profile:*

Step 1. Verify the appropriate license is installed. Run `show licenses`.

Step 2. To change the system profile to:

- an Ethernet system profile, run `system profile eth-single-switch`.
- an InfiniBand system profile, run `system profile ib-single-switch`.
- InfiniBand system profile without adaptive system profile, run `ib-no-adaptive-routing-single-switch`.
- IB router system profile, run `system profile ib num-of-swids`.
- a VPI system profile, run `system profile vpi-single-switch`.

Upon approval, the configuration is deleted and the switch is rebooted with the new profile.

Step 3. To verify the system profile, run `show system profile`.



In case the system profile is `vpi-single-switch`, it is possible to change the link protocol from InfiniBand to Ethernet or vice versa using “port type” command.

4.5.1 Commands

system profile

system profile {eth-single-switch | ib-single-switch | ib-no-adaptive-routing-single-switch | ib num-of-swids <swid-num> [ib-router] [adaptive-routing] | vpi-single-switch} [force]

Sets the profile of the system to either InfiniBand, IB Router, Ethernet, or VPI.

Syntax Description	ib-single-switch	Enables InfiniBand switch profile All network interfaces link protocol set to InfiniBand
	ib-no-adaptive-routing-single-switch	Enables InfiniBand switch profile without adaptive routing capabilities All network interfaces link protocol set to InfiniBand with disabled adaptive routing
	ib num-of-swids	Enables IB Router Multiple switch IDs are configurable • adaptive routing – enables adaptive routing • ib-router – enables IB router
	eth-single-switch	Enables Ethernet switch profile All network interfaces link protocol set to Ethernet
	vpi-single-switch	Enables VPI switch profile Some ports can be defined as Ethernet while some other as InfiniBand
	force	Force operation, without the need for user confirmation.
Default	The default system profile depends on the system: SX6xxx systems have “ib-single-switch” as default SX1xxx systems have “eth-single-switch” as default SX6036G system has “vpi-single-switch” as default SB7780 system has “IB Router” and 2 SWIDs as default	
Configuration Mode	Config	
History	3.1.0000	First version
	3.2.1100	Added “vpi-single-switch” option
	3.3.4100	Added SX6036G
	3.3.4302	Added system profile ib-no-adaptive-routing-single-switch

	3.6.1002	Added system profile “ib num-of-swids”
Role	admin	
Example	<pre>switch (config) # system profile eth-single-switch switch (config) #</pre>	
Related Commands	<pre>port type show system profile show ports type</pre>	
Notes	• This command requires a license. Refer to “Licenses” section in the <i>MLNX-OS SwitchX User Manual</i> • This command requires approval because reboot is performed and all configuration is removed • This command deletes all switch configuration (keeping IP connectivity) and resets the system • System profile “ib-no-adaptive-routing-single-switch profile” is the default profile for InfiniBand switches • Adaptive routing is not available on system profile “vpi-single-switch” • Refer to the ‘port type’ command in order to change the link protocol • IB router and adaptive routing are enabled only if specified but cannot be enabled at the same time	

4.6 Software Management



To interoperate with Switch-IB™ 2 based switch systems, switch systems must at least be installed with MLNX-OS version 3.6.1002.

4.6.1 Upgrading MLNX-OS Software



When upgrading from a software version older than 3.2.0100 to software version 3.3.0000 or higher, the upgrade procedure must be done in two steps. First update the software to 3.2.0300-100 (for InfiniBand platforms) or 3.2.0506 (for Ethernet platforms), then update to the desired software version.



Upgrading director switch systems can take up to 30 minutes during which time the system is indisposed.



The system being upgraded becomes indisposed throughout the upgrade procedure.



The upgrade procedure burns the software image as well as the firmware should there be a need.



If running a system with dual management cards, refer to [Section 4.6.2, “Upgrading MLNX-OS Software on Director Switches,”](#) on page 246.



To upgrade the MLNX-OS version on a gateway, SM, or MLAG cluster, please refer to [Section 4.6.3, “Upgrading MLNX-OS HA Groups,”](#) on page 247.



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.

To upgrade MLNX-OS software on your system, perform the following steps:

Step 1. Change to Config mode.

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

- Step 2.** Obtain the previously available image (.img file). You *must* delete this image in the next step to make room for fetching the new image.

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

Partition 1:
SX_PPC_M460EX 3.3.3130 2013-03-20 21:32:25 ppc

Partition 2:
SX_PPC_M460EX 3.3.3130 2013-03-20 21:32:25 ppc

Images available to be installed:

image-PPC_M460EX-SX_3.3.3256.img
SX_PPC_M460EX 3.3.3256 2013-03-20 21:32:25 ppc

Serve image files via HTTP/HTTPS: no

No image install currently in progress.

Boot manager password is set.

No image install currently in progress.

Require trusted signature in image being installed: yes (default)
switch (config) #
```

- Step 3.** Delete the old image (if one exists) that is 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 image-PPC_M460EX-3.0.1224.img
switch (config) #
```



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

- Step 4.** Fetch the new software image.

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

Step 5. Display the available images.



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
image boot location.
```

```
switch (config) # show images
Installed images:
  Partition 1:
    SX <old ver> 2013-04-28 16:02:50

  Partition 2:
    SX <new ver> 2013-04-28 16:52:50

Images available to be installed:
  new_image.img
  SX <new ver> 2013-04-28 16:52:50

Serve image files via HTTP/HTTPS: no

No image install currently in progress.

Boot manager password is set.

No image install currently in progress.

Require trusted signature in image being installed: yes (default)
switch (config) #
```

Step 6. Install the new image.

```
switch (config) # image install <image_name>
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) #
```



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

Step 7. Have the new image activate during the next boot. Run:

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

Step 8. Run `show images` to review your images. Run:

```
switch (config) # show images
Images available to be installed:
  new_image.img
  SX <new ver> 2011-04-28 16:52:50

Installed images:
  Partition 1:
  SX <old ver> 2011-04-28 16:02:50

  Partition 2:
  SX <new ver> 2011-04-28 16:52:50

Last boot partition: 1
Next boot partition: 2

No boot manager password is set.
switch (config) #
```

Step 9. Save current configuration. Run:

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

Step 10. Reboot the switch to run the new image. Run:

```
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 firm-ware version.



On SX65xx systems with dual management, the software must be upgraded on both the master and the slave modules.



In order to upgrade the system on dual management system refer to [Section 4.6.1, “Upgrading MLNX-OS Software,”](#) on page 242.



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

4.6.2 Upgrading MLNX-OS Software on Director Switches



Director switches feature dual management modules.

Step 1. Identify the chassis HA master. Run:

```
show chassis ha
```

Step 2. Upgrade the chassis master according to steps 1-8 in section [Section 4.6.1, on page 242](#). Please DO NOT reboot.

Step 3. Upgrade the second management module according to steps 1-8 in section [Section 4.6.1, on page 242](#). Please DO NOT reboot.

Step 4. Reset the slave management module. In the master management module, run:

```
chassis ha reset other
```

Step 5. After invoking the command above, please reboot the master management immediately. Run:

On SX65xx switch systems, run:

```
reload
```

On CS75x0 switch systems, run:

```
reload force immediate
```



An alternative for [Step 4](#) and [Step 5](#) is to power cycle the system.

Step 6. Check that 'reset count' equals 0 or 1. Run:

```
show chassis ha
```

If the reset count is not equal to either 0 or 1, power cycle the system.

Step 7. Verify all the systems are back online as members of the IB subnet ID. Run:

```
show ib smnodes {brief}
```



Using a director switch with different software versions on its two management boards is not supported.

When replacing a management board the software running on the replacement board must be aligned with the version of the software running on the other management board.

4.6.3 Upgrading MLNX-OS HA Groups

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



For Proxy-ARP HA, the procedure below is valid from MLNX-OS v3.4.1120 and later.

➤ **To upgrade MLNX-OS version without affecting an HA group:**

Step 1. Identify the HA group master.

for IB HA. Run:

```
switch (config) # show ib ha
Global HA state
=====
IB Subnet HA name:subnet4
HA IP address: 192.168.10.43/24
Active HA nodes: 2
ID              State Role      IP              SM Priority
-----
switch          standalone 192.168.10.42   disabled
switch          master    192.168.10.18   disabled
```

for MLAG. Run:

```
switch (config)# show mlag-vip
MLAG VIP
=====
MLAG group name: my-mlag-group
MLAG VIP address: 1.1.1.1/30
Active nodes: 2
Hostname        VIP-State      IP Address
-----
SwitchA         master         10.10.10.1
SwitchB         standby        10.10.10.2
```

For Gateway HA. Run:

```
GatewayA [my-group: master] (config) # show proxy-arp ha
Load balancing: ib-base-ip
Number of Proxy-ARP interfaces: 1

Proxy Arp VIP:
=====
Proxy-arp group name: my-group
HA VIP address: 10.10.10.10/24
Active nodes: 2
Hostname  State  IP Address
-----
GatewayA  master 10.10.10.11
GatewayB  standby 10.10.10.12
```

- Step 2.** Upgrade standby nodes in the HA group according to steps 1-10 in section [Section 4.6.1, on page 242](#).
- Step 3.** Wait until all standby nodes have rejoined the group.
- Step 4.** Upgrade the master node in the HA group according to steps 1-10 in section [Section 4.6.1, on page 242](#).

4.6.4 Deleting Unused Images

➤ *To delete unused images:*

- Step 1.** Enter Config mode. Run:

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

- Step 2.** Get a list of the unused images. Run

```
switch (config) # show images
Images available to be installed:
  image-PPC_M460EX-3.1.1224.img
  SX-OS_PPC_M460EX 3.1.1224 2011-04-28 12:29:48 ppc
Installed images:
Partition 1:
SX-OS_PPC_M460EX 3.1.0000-dev-HA 2011-04-10 12:02:49 ppc
Partition 2:
SX-OS_PPC_M460EX 3.1.0000-dev-HA 2011-04-10 12:02:49 ppc

Last boot partition: 1
Next boot partition: 1
Boot manager password is set.
No image install currently in progress.
Require trusted signature in image being installed: yes
switch (config) #
```

- Step 3.** Delete the unused images. Run:

```
switch config) # image delete image-PPC_M460EX-3.1.1224.img
switch (config) #
```



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

4.6.5 Downgrading MLNX-OS Software

IMPORTANT NOTE

If in possession of an MSX1xxx or MSX6xxx switch system housing a SwitchX®-2 IC, then the oldest MLNX-OS version to which you may downgrade is 3.3.5006; otherwise, the switch system will malfunction.

To find out whether your system is based on SwitchX®-2, please run the command “show inventory” and make sure that the “Asic Rev” column of the MGMT indicates “2”.

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

Step 1. Log into your switch via the CLI using the console port.

Step 2. Backup your configuration according to the following steps:

1. Change to Config mode. Run:

```
switch-112094 [standalone: master] > enable
switch-112094 [standalone: master] # configure terminal
switch-112094 [standalone: master] (config) #
```

2. Disable paging of CLI output. Run:

```
switch-112094 [standalone: master] (config) # no cli default paging enable
```

3. Display commands to recreate current running configuration. Run:

```
switch-112094 [standalone: master] (config) # show running-config
```

4. Copy the output to a text file.

4.6.5.1 Downloading Image

Step 1. Log into your system to obtain its product number. Run:

```
switch-112094 [standalone: master] (config) # show inventory
```

Step 2. Log into MyMellanox at <https://mymellanox.force.com/support/SupportLogin> and download the relevant MLNX-OS version to your system type.

Step 3. Log into the switch via the CLI using the console port.

Step 4. Change to Config mode. Run:

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

Step 5. Delete all previous images from the Images available to be installed prior to fetching the new image. Run:

```
switch (config) # image delete image-EFM_PPC_M405EX-ppc-m405ex 20090531-190132.img
```

Step 6. Fetch the requested software image. Run:

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

4.6.5.2 Downgrading Image



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

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

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

Step 3. Show all image files on the system. Run:

```
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.
switch (config) #
```

Step 4. Install the MLNX-OS image. Run:

```
switch (config) # image install <image_name>
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) #
```

Step 5. Show all image files on the system. Run:

```
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.
switch (config) #
```

Step 6. Set the boot location to be the other (next) partition. Run:

```
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.



In case you are downloading 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

Step 7. Reload the switch. Run:

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

4.6.5.3 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.



IMPORTANT NOTE

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 than the backup file

Note that the configuration file becomes empty if the switch 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:

Step 1. Run the command:

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

Step 2. Set the boot partition. Run:

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

Step 3. Save the configuration. Run:

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

Step 4. Reload the system. Run:

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

4.6.6 Upgrading System Firmware

Each MLNX-OS software package version has a default switch firmware version. When you update the MLNX-OS software to a new version, an automatic firmware update process will be attempted by MLNX-OS. This process is described below.

4.6.6.1 After Updating MLNX-OS Software

Upon rebooting your switch system after updating the MLNX-OS software, MLNX-OS compares its default firmware version with the currently programmed firmware versions on all the switch modules (leafs and spines on director-class switches, or simply the switch card on edge switch systems).

If one or more of the switch modules is programmed with a firmware version other than the default version, then MLNX-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 MLNX-OS 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 Mellanox Technologies field application engineer.

4.6.6.2 After Inserting a Switch Spine or Leaf



This section is applicable to director-class switch systems only.

If you insert a switch spine or leaf with a firmware version other than the default version of MLNX-OS, an automatic firmware update process will take place immediately to the inserted module *only*.



The firmware update may take a few minutes. It is recommended not to run any commands until the firmware update completes.



During firmware upgrade internal link status (up/down) notifications may be sent.

To verify that the firmware update was successful, run the command “show asic-version” (can be run in any mode). Check that the firmware version of the inserted switch spine or leaf has the default firmware version.



If you detect a firmware version mismatch for the newly inserted module, please contact your assigned Mellanox Technologies field application engineer.

4.6.6.3 Importing Firmware and Changing the Default Firmware

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



From version 3.3.4400 and above, the firmware update file format has been changed to mfa format. TGZ format is no longer supported.

4.6.6.3.1 Default Firmware Change on Standalone Systems

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

```
switch (config) # image fetch scp://root@1.1.1.1:/tmp/fw-SX-rel-9_2_6440-FIT.mfa
Password (if required): *****
100.0%
[#####]
switch (config) # image default-chip-fw fw-SX-rel-9_2_6440-FIT.mfa
Installing default firmware image. Please wait...
Default Firmware 9.2.6440 updated. Please save configuration and reboot for new FW to
take effect.
switch (config) #
```

Step 2. Save the configuration. Run:

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

Step 3. Reboot the system to enable auto update.

4.6.6.3.2 Default Firmware Change Dual Management Systems

This flow should be implemented on both Orca managements in parallel.

Step 1. Import the firmware image (.mfa file) on both management modules.

```
switch (config) # image fetch scp://username:password@10.7.34.115//my_directory/fw-SX-
rel-9_1_6470-FIT.mfa
100.0%
[#####]
```

Step 2. Change default firmware on the management modules using the command `image default-chip-fw`.

Step 3. Verify that both master and slave have successfully installed the new firmware. The following message should be displayed:

```
Default firmware <fw> updated. Please save configuration and reboot for new FW to take
effect.
```

Step 4. Run configuration write on both management modules.

Step 5. Run `chassis ha reset other` on master only.

Step 6. Run reload on master only.

4.6.7 Image Maintenance via Mellanox ONIE



Supported only on MSX1710-BS2F2O, and Mellanox Spectrum™ based switch systems.

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.

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

➤ **To change the switch system's OS:**

Step 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
```

Step 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 ?
```

Step 3. Type YES to continue.

Since MLNX-OS 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
...
...
```

Step 4. In order to manually insert an install URL, press Enter and insert the command “install_url <http> / <tftp> <url> <image name .bin>”. For example:

```
install_url http://<ip_address>/sx_mlnx_os/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.

Step 5. Wait for the installation to end and reboot this switch to boot into the OS.

```
ONIE:/ # install_url http://<ip_address>//sx_mlnx_os/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/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/sx_mlnx_os-3.5.1000-21/
X86_64/X86_64-3.5.1000-21-installer.bin
```


4.6.8 Commands

This chapter displays all the relevant commands used to manage the system software image.

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	enable/config	
History	3.1.0000	
Role	admin	
Example	switch (config) # image boot location 2 switch (config) #	
Related Commands	show images	
Notes		

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
Role	admin
Example	switch (config) # boot next fallback-reboot enable switch (config) #
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 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 <i>has</i> 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.

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	
Role	admin	
Example	switch (config) # boot system location 2 switch (config) #	
Related Commands	show images	
Notes		

image default-chip-fw

image default-chip-fw <file name>

Sets the default firmware package to be installed.

Syntax Description	filename	Specifies the firmware filename.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # image default-chip-fw fw-SX-rel-9_2_6440-FIT.mfa	
Related Commands	show asic-version show images	
Notes		

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	
Role	admin	
Example	switch (config) # image delete image-MLXNX-OS-201140526-010145.img switch (config) #	
Related Commands	show images	
Notes		

image fetch

image fetch <URL> [<filename>]

Downloads an image from the specified URL or via SCP.

Syntax Description	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	
Role	admin	
Example	<pre>switch (config) # image fetch scp://<username>@192.168.10.125/var/www/ html/<image_name> Password ***** 100.0%[#####] switch (config) #</pre> <u>Other options:</u> <pre>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</pre>	

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 section “Upgrading MLNX-OS SX Software,” in the <i>Mellanox SwitchX® User Manual</i> for a full upgrade example
-

image install

image install <image filename> [location <location ID>] | [progress <prog-options>] [verify <ver-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
	ver-options	“check-sig” requires an image to have either a valid signature or no signature “ignore-sig” allows unsigned or invalidly signed images to be installed “require-sig” requires from the installed image to have a valid signature. If a valid signature is not found on the image, the image cannot be installed.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # image install SX_PPC_M460EX 3.0.0000-dev-HA 2012-01-22 08:47:59 ppc 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) #</pre>	

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
-
-

image move

image move <src image name> <dest image name>

Renames the specified image file.

Syntax Description	src image name	Specifies the old image name.
	dest image name	Specifies the new image name.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # image move image1.img image2.img switch (config) #	
Related Commands	show images	
Notes		

image options

image options {require-sig | serve all}
no image options {require-sig | serve all}

Configures options and defaults for image usage.
The no form of the command disables options and defaults for image usage.

Syntax Description	require-sig	Requires images to be signed by a trusted signature
	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	
Role	admin	
Example	switch (config) # image options require-sig	
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>	

show bootvar

show bootvar

Displays the installed system images and the boot parameters.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre> switch (config) # show bootvar Installed images: Partition 1: SX_PPC_M460EX 3.0.0000-dev-HA 2012-01-22 08:47:59 ppc Last dobincp: 2012/01/23 14:54:23 Partition 2: SX_PPC_M460EX 3.0.0000-dev-HA 2012-01-18 09:52:41 ppc Last dobincp: 2012/01/19 16:48:23 Last boot partition: 1 Next boot partition: 1 Boot manager password is set. No image install currently in progress. Image signing: trusted signature always required Admin require signed images: yes Settings for next boot only: Fallback reboot on configuration failure: yes (default) switch (config) # </pre>
Related Commands	N/A
Notes	

show images

show image

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
Role	admin
Example	<pre>switch (config) # show images Images available to be installed: image-SX_PPC_M460EX-ppc-m460ex-20120122-084759.img SX_PPC_M460EX 3.0.0000-dev-HA 2012-01-22 08:47:59 ppc Installed images: Partition 1: SX_PPC_M460EX 3.0.0000-dev-HA 2012-01-22 08:47:59 ppc Last dobincp: 2012/01/23 14:54:23 Partition 2: SX_PPC_M460EX 3.0.0000-dev-HA 2012-01-18 09:52:41 ppc Last dobincp: 2012/01/19 16:48:23 Last boot partition: 1 Next boot partition: 1 Boot manager password is set. No image install currently in progress. Image signing: trusted signature always required Admin require signed images: yes Settings for next boot only: Fallback reboot on configuration failure: yes (default) switch (config) #</pre>
Related Commands	N/A
Notes	

4.7 Configuration Management

4.7.1 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) #
```

4.7.2 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 [standalone: master] >
switch [standalone: master] > enable
switch [standalone: master] # configure terminal
switch [standalone: master] (config) # configuration switch-to myconfig
switch [standalone: master] (config) #
```

On director switch systems with dual management modules, load the configuration file according to the following:

- Step 1.** Power cycle the system.
- Step 2.** Load the configuration on the top CPU that serves as the chassis master according to the procedure described above.



If the configuration file is loaded on a different CPU than the SM HA master (SM HA master that serves the VIP), the SM configuration is overwritten.

4.7.3 Restoring Factory Default Configuration

In cases where the system configuration becomes corrupted it is suggested to restore the factory default configuration.

➤ **To restore factory default configuration on a single management module system:**

Step 1. Run the command `reset factory [reboot] [keep-basic] [keep-all-config]:`

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

➤ **To restore factory default configuration on a dual management module system:**

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

Step 1. Connect to a remote console/serial connection.

Step 2. Remove the slave management module.

Step 3. Run the command `reset factory [keep-basic] [keep-all-config]:`

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

Please wait for reboot to complete before moving to the next step.

Step 4. Log in as “admin” and start running the Mellanox Configuration Wizard.

Step 5. Insert the slave management module.

Step 6. Remove the master management module.



A takeover will occur changing the Slave management module role to Master.

Step 7. Repeat Step 3 on the new Master management module.

Step 8. Insert the other management module. No takeover will occur at this stage.

Step 9. Power cycle the system.

4.7.4 Managing Configuration Files

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

4.7.4.1 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 on the switch.

➤ **To create a new BIN configuration file:**

```
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 from a switch to an external file server:**

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

- **To fetch a BIN configuration file:**

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

- **To see the available configuration files:**

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

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

- **To load a BIN configuration file:**

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



Applying a new BIN configuration file changes the whole switch's configuration and requires system reboot which can be preformed using the command `reload`.



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.

4.7.4.2 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:**

```
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:***

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



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 switch's configuration by applying a specific configuration file (following the procedure in [Section 4.7.4.1](#)) or by resetting the switch back to factory default.

➤ ***To upload a text-based configuration file from a switch to an external file server***

```
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***

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

➤ ***To apply a text-based configuration file:***

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



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

4.7.5 Commands

4.7.5.1 File System

debug generate dump

debug generate dump

Generates a debug dump.

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

file debug-dump

file debug-dump {delete {<filename> | latest} | email {<filename> | latest} | upload {{<filename> | latest} <URL>}}

Manipulates debug dump files.

Syntax Description	delete {<filename> latest}	Deletes a debug dump file.
	email {<filename> latest}	Emails a debug dump file to pre-configured recipients for “informational events”, regardless of whether they have requested to receive “detailed” notifications or not.
	upload {{<filename> latest} <URL>}}	Uploads a debug dump file to a remote host. The URL to the remote host: HTTP, HTTPS, FTP, TFTP, SCP and SFTP are supported. Example: scp://username[:password]@hostname/path/filename.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	Initial release
	3.3.4000	Added “latest” parameter
Role	admin	
Example	switch (config) # file debug-dump email sysdump-switch-112104-20114052-091707.tgz switch (config) #	
Related Commands	show files debug-dump	
Notes		

file stats

file stats {delete <filename> | move {<source filename> | <destination filename>}
| upload <filename> <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. URL - HTTP, HTTPS, FTP, TFTP, SCP and SFTP are supported. Example: scp://username[:password]@hostname/path/filename.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # file stats move memory-1.csv memory-2.csv switch (config) #	
Related Commands	show files stats show files stats <filename>	
Notes		

file tcpdump

file tcpdump {delete <filename> | upload <filename> <URL>}

Manipulates tcpdump output files.

Syntax Description	delete <filename>	Deletes the specified tcpdump output file.
	upload <filename> <URL>	Uploads the specified tcpdump output file to the specified URL. URL - HTTP, HTTPS, FTP, TFTP, SCP and SFTP are supported. Example: scp://username[:password]@hostname/path/filename.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # file tcpdump delete my-tcpdump-file.txt switch (config) #	
Related Commands	show files stats tcpdump	
Notes		

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.
	noconfirm	Reboots the system without asking about unsaved changes.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # reload Configuration has been modified; save first? [yes] yes Configuration changes saved. ... switch (config) #</pre>	
Related Commands	reset factory	
Notes	BBU discharge must be disabled before any planned shutdown of the switch	

reset factory

reset factory [keep-all-config | keep-basic | keep-virt-vols | only-config] [halt]

Clears the system and resets it entirely to its factory state.

Syntax Description	keep-all-cofig	Preserves all configuration files including licenses. Removes the logs, stats, images, snapshots, history, 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-virt-vols	Preserve all virtual disk volumes
	only-config	Removes configuration files only. The logs, stats, images, snapshots, history, and known hosts are preserved.
	halt	The system is halted after this process completes
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
Role	admin	
Example	<pre>switch (config) # reset factory Warning - confirming will cause system reboot. Type 'YES' to confirm reset: YES Resetting and rebooting the system -- please wait... ...</pre>	

Related Commands

reload

Notes

- Effects of parameter “keep-all-cofig”: 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 – unchanged
 - Confirming the command causes system reboot
-
-

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	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # show files debug-dump sysdump-switch-112104-20114052-091707.tgz System information: Hostname: switch-112104 Version: SX_PPC 3.1.0000 2011-05-25 13:59:00 ppc Date: 2012-01-26 09:17:07 Uptime: 0d 18h 47m 48s ===== Output of 'uname -a': Linux switch-112104 2.6.27-MELLANOXuni-m405ex SX_PPC 3.1.0000 #1 2012-01-25 13:59:00 ppc ppc ppc GNU/Linux ===== switch (config) #</pre>	
Related Commands	file debug-dump	
Notes		

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	Config
History	3.1.0000
Role	admin
Example	switch (config) # show files stats memory-201140524-111745.csv switch (config) #
Related Commands	file stats
Notes	

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	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # show files system Statistics for /config filesystem: Bytes Total 100 MB Bytes Used 3 MB Bytes Free 97 MB Bytes Percent Free 97% Bytes Available 97 MB Inodes Total 0 Inodes Used 0 Inodes Free 0 Inodes Percent Free 0% Statistics for /var filesystem: Bytes Total 860 MB Bytes Used 209 MB Bytes Free 651 MB Bytes Percent Free 75% Bytes Available 651 MB Inodes Total 0 Inodes Used 0 Inodes Free 0 Inodes Percent Free 0% switch (config) #</pre>	
Related Commands	N/A	
Notes		

show files tcpdump

show files tcpdump

Displays a list of statistics report files.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show files stats test dump3 switch (config) #
Related Commands	file tcpdump tcpdump
Notes	

4.7.5.2 Configuration Files

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	
Role	admin	
Example	switch (config) # configuration audit max-changes 100 switch (config) # show configuration audit Maximum number of changes to log: 100 switch (config) #	

Related Commands	show configuration

Notes	N/A

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	
Role	admin	
Example	switch (config) # configuration copy initial.bak example switch (config) #	
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.	

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	
Role	admin	
Example	<pre>switch (config) # show configuration files example initial initial.bak initial.prev switch (config) # configuration delete example switch (config) # show configuration files initial initial.bak initial.prev switch (config) #</pre>	
Related Commands	show configuration	
Notes	- This command does not affect the current running configuration - The active configuration file may not be deleted	

configuration fetch

configuration fetch <URL> [<name>]

Downloads a configuration file from a remote host.

Syntax Description	URL	HTTP, HTTPS, FTP, TFTP, SCP and SFTP are supported. Example: scp://username[:password]@hostname/path/filename.
	name	The configuration file name.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # configuration fetch scp://root:password@ 192.168.10.125/tmp/conf1 switch (config) #</pre>	
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”	

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
Role	admin
Example	<pre>switch (config) # configuration jump-start Mellanox 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. switch (config) #</pre>
Related Commands	N/A
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 chapter “Initializing the Switch for the First Time” in the Mellanox <i>MLNX-OS SwitchX User Manual</i>

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	
Role	admin	
Example	switch (config) # configuration merge new-config-file switch (config) #	
Related Commands		
Notes	• No configuration files are modified during this process • The configuration name must be a non-active configuration file	

configuration move

configuration move <source name> <dest name>

Moves a configuration file.

Syntax Description	source name	Old name of file to move.
	dest name	New name for moved file.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>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 switch (config) #</pre>	
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	

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	
Role	admin	
Example	<pre>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 switch (config) #</pre>	
Related Commands	show configuration	
Notes		

configuration revert

configuration revert {factory [keep-basic | keep-connect] | saved}

Reverts the system configuration to a previous state.

Syntax Description	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).
	saved	Reverts running configuration to last saved configuration.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # configuration revert saved switch (config) #	
Related Commands	show configuration	
Notes	This command is only available when working with an InfiniBand profile This command is not available on IB multi-SWID system profile	

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 the switch
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
	3.6.1002	Added “no-reboot” option
Role	admin	
Example	<pre>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 switch (config) #</pre>	
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 switch is rebooted immediately.	

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	First version
	3.2.3000	Updated command
Role	admin	
Example	switch (config) # configuration fetch text scp://username[:password]@hostname/path/filename	
Related Commands	N/A	
Notes		

configuration text file

configuration text file <filename> {**apply** [**fail-continue**] [**verbose**] | **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.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # configuration text file my-config-file delete switch (config) #	
Related Commands	show configuration files	
Notes		

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	
Role	admin	
Example	<pre>switch (config) # configuration text generate file initial.prev save example switch (config) # show configuration files initial (active) initial.prev initial.bak Active configuration: initial Unsaved changes: yes switch (config) #</pre>	
Related Commands	show configuration files	
Notes		

configuration upload

configuration upload {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.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # configuration upload active scp://root:password@ 192.168.10.125/tmp/conf1 switch (config) #</pre>	
Related Commands	N/A	
Notes	No configuration file may have the name “active”.	

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	
Role	admin	
Example	switch (config) # configuration write switch (config) #	
Related Commands	write	
Notes		

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	
Role	admin	
Example	<pre> switch (config) # write terminal ## ## Running database "initial" ## Generated at 20114/05/27 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 switch (config) # </pre>	

Related Commands show running-config
 configuration write

Notes

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 “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
Role	monitor/admin

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
switch (config) #
```

Related Commands

Notes

show running-config

show running-config [expanded | protocol <protocol>]

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
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
	3.3.4402	Removed “full” parameter
	3.6.2002	Updated Example and added parameters
Role	monitor/admin	

Example

```
switch (config) # show running-config
##
## Running database "initial"
## Generated at 2016/08/03 17:28:18 +0000
## Hostname: tarantula-9
##
##
## Running-config temporary prefix mode setting
##
no cli default prefix-modes enable

##
## Chassis configuration
##
no fae fw-auto-update enable

##
## MLAG protocol
##
    protocol mlag

##
## Interface Ethernet configuration
##
    interface mlag-port-channel 1-49
    interface mlag-port-channel 53-56
    interface port-channel 1
    interface ethernet 1/1-1/43 mtu 9216 force
    interface ethernet 1/49-1/56 mtu 9216 force
    interface mlag-port-channel 1-42 mtu 9216 force
    interface mlag-port-channel 49 mtu 9216 force
    interface mlag-port-channel 53 mtu 9216 force
...
switch (config) #
```

Related Commands

Notes

4.8 Logging

4.8.1 Monitor

➤ *To print logging events to the terminal:*

Set the modules or events you wish to print to the terminal. For example, run:

```
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
switch (config) #
```

To see a list of the events, refer to [Table 28, “Supported Event Notifications and MIB Mapping,”](#) on page 351.

4.8.2 Remote Logging

➤ *To configure remote syslog to send syslog messages to a remote syslog server:*

Step 1. Enter Config mode. Run:

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

Step 2. Set remote syslog server. Run

```
switch (config) # logging <IP address>
```

Step 3. (Optional) Set the destination port of the remote host. Run:

```
switch (config) # logging <IP address> port <port>
```

Step 4. Set the minimum severity of the log level to info. Run:

```
switch (config) # logging <IP address> trap info
```

Step 5. Override the log levels on a per-class basis. Run:

```
switch (config) # logging <IP address> trap override class <class name> priority
<level>
```

4.8.3 Commands

logging <syslog IP address> port

logging <syslog IP address> port <destination-port>
no logging <syslog IP address> 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
Default	514 (UDP)	
Configuration Mode	Config	
History	3.6.2002	
Role	admin	
Example	switch (config) # logging 10.0.0.1 port 105	
Related Commands	logging <syslog IP address> trap	
Notes		

logging <syslog IP address> trap

logging <syslog IP address> [trap {<log-level> | override class <class> priority <log-level>}]

no logging <syslog IP address> [trap {<log-level> | override class <class> priority <log-level>}]

Enables (by setting the IP address) 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 IP address	IPv4 address of the remote syslog server.
	log-level	• alert - alert notification, action must be taken immediately • crit - critical condition • debug - debug level messages • emerg - system is unusable (emergency) • err - error condition • info - informational condition • none - disables the logging locally and remotely • notice - normal, but significant condition • warning - warning condition
	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, MLNX-OS 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
	log-level	• alert - alert notification, action must be taken immediately • crit - critical condition • debug - debug level messages • emerg - system is unusable (emergency) • err - error condition • info - informational condition • none - disables the logging locally and remotely • notice - normal, but significant condition • warning - warning condition
Default	Remote logging is disabled	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	

Example

```
switch (config) # logging local info
switch (config) # show logging
Local logging level: info
Default remote logging level: notice
No remote syslog servers configured.
Allow receiving of messages from remote hosts: no
Number of archived log files to keep: 10
Log rotation size threshold: 5.000% of partition (43 megabytes)
Log format: standard
Subsecond timestamp field: disabled
Levels at which messages are logged:
  CLI commands: notice
  Audit messages: notice
switch (config) #
```

Related Commands

```
show logging
logging local override
logging <syslog IP address> port
```

Notes

logging debug-files

logging debug-files {delete {current | oldest} | rotation {criteria | force | max-num} | update {<number> | current} | upload <log-file> <upload URL>}

Configures settings for debug log files.

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	HTTP, HTTPS, FTP, TFTP, SCP and SFTP are supported (e.g.: scp://username[:password]@host-name/path/filename)
Default	N/A	
Configuration Mode	Config	
History	3.3.4150	
Role	admin	

Example

```
switch (config) # logging debug-files delete current  
switch (config) #
```

Related Commands

Notes

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, MLNX-OS 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	• alert - alert notification, action must be taken immediately • crit - critical condition • debug - debug level messages • emerg - system is unusable (emergency) • err - error condition • info - informational condition • none - disables the logging locally and remotely • notice - normal, but significant condition • warning - warning condition
Default	Override is disabled.
Configuration Mode	Config

History	3.1.0000	
	3.3.4150	Added debug-module class Changed iss-modules with protocol-stack
Role	admin	
Example	<pre>switch (config) # logging local override class mgmt-front priority warning switch (config) # show logging Local logging level: info Override for class mgmt-front: warning Default remote logging level: notice No remote syslog servers configured. Allow receiving of messages from remote hosts: no Number of archived log files to keep: 10 Log rotation size threshold: 5.000% of partition (43 megabytes) Log format: standard Subsecond timestamp field: disabled Levels at which messages are logged: CLI commands: notice Audit messages: notice switch (config) #</pre>	
Related Commands	<pre>show logging logging local</pre>	
Notes		

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	
Role	admin	
Example	<pre>switch (config) # logging fields seconds enable switch (config) # logging fields seconds whole-digits 1 switch (config) # show logging Local logging level: info Override for class mgmt-front: warning Default remote logging level: notice No remote syslog servers configured. Allow receiving of messages from remote hosts: no Number of archived log files to keep: 10 Log rotation size threshold: 5.000% of partition (43 megabytes) Log format: standard Subsecond timestamp field: enabled Subsecond timestamp precision: 1 whole digit; 3 fractional digits Levels at which messages are logged: CLI commands: notice Audit messages: notice switch (config) #</pre>	

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.

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	
Role	admin	
Example	switch (config) # logging files delete current switch (config) #	
Related Commands	show logging show log files	
Notes		

logging files rotation

logging files rotation {criteria { frequency <freq> | size <size-mb>| size-pct <size-percentage>} | force | max-number <number-of-files>}

Sets the rotation criteria of the logging files.

Syntax Description	freq	Sets rotation criteria according to time. Possible options are: • Daily • Weekly • Monthly
	size-mb	Sets rotation criteria according to size in mega bytes. The range is 1-9999.
	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	
Role	admin	

Example

```
switch (config) # logging files rotation criteria size-pct 6
switch (config) # show logging
Local logging level: info
  Override for class mgmt-front: warning
Default remote logging level: notice
No remote syslog servers configured.
Allow receiving of messages from remote hosts: no
Number of archived log files to keep: 10
Log rotation size threshold: 6.000% of partition (51.60 megabytes)
Log format: standard
Subsecond timestamp field: enabled
Subsecond timestamp precision: 1 whole digit; 3 fractional digits
Levels at which messages are logged:
  CLI commands: info
  Audit messages: notice
switch (config)
```

Related Commands

```
show logging
show log files
```

Notes

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. FTP, TFTP, SCP, and SFTP are supported. 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	
Role	admin	
Example	switch (config) # logging files upload 1 scp://admin@scpserver	
Related Commands	show logging show log files	
Notes		

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.
Default	standard	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # logging format standard switch (config) # show logging Local logging level: info Default remote logging level: notice No remote syslog servers configured. Allow receiving of messages from remote hosts: yes Number of archived log files to keep: 10 Log rotation size threshold: 5.000% of partition (43 megabytes) Log format: standard Subsecond timestamp field: disabled Levels at which messages are logged: CLI commands: notice Audit messages: notice switch (config) #</pre>	
Related Commands	show logging	
Notes		

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	• alert - alert notification, action must be taken immediately • crit - critical condition • debug - debug level messages • emerg - system is unusable (emergency) • err - error condition • info - informational condition • none - disables the logging locally and remotely • notice - normal, but significant condition • warning - warning condition
Default	CLI commands and audit message are set to notice logging level	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre> switch (config) # logging level cli commands info switch (config) # show logging Local logging level: info Override for class mgmt-front: warning Default remote logging level: notice No remote syslog servers configured. Allow receiving of messages from remote hosts: no Number of archived log files to keep: 10 Log rotation size threshold: 5.000% of partition (43 megabytes) Log format: standard Subsecond timestamp field: enabled Subsecond timestamp precision: 1 whole digit; 3 fractional digits Levels at which messages are logged: CLI commands: info Audit messages: notice switch (config) # </pre>	

Related Commands show logging

Notes

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	
Role	admin	
Example	<pre>switch (config) # logging monitor events notice switch (config) #</pre>	
Related Commands		
Notes		

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
Role	admin
Example	<pre> switch (config) # logging receive switch (config) # show logging Local logging level: info Default remote logging level: notice No remote syslog servers configured. Allow receiving of messages from remote hosts: yes Number of archived log files to keep: 10 Log rotation size threshold: 5.000% of partition (43 megabytes) Log format: standard Subsecond timestamp field: disabled Levels at which messages are logged: CLI commands: notice Audit messages: notice switch (config) # </pre>
Related Commands	<pre> show logging logging local logging local override </pre>
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

logging trap

logging trap <log-level>
no logging trap

Configures the minimum severity of log messages sent to syslog servers.
 The no form of the command disables sending event log messages to syslog servers.

Syntax Description	log-level	The minimum severity level for all configured syslog servers: • none – disable logging • emerg – emergency: system is unusable • alert – action must be taken immediately • crit – critical conditions • err – error conditions • warning – warning conditions • notice – normal but significant condition • info – informational messages • debug – debug-level messages
Default	Receiving logging is disabled	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # logging trap info switch (config) #</pre>	
Related Commands		
Notes		

show logging

show logging

Displays the logging configurations.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.1.0000
Role	admin
Example	<pre>switch (config) # show logging Local logging level: info Override for class mgmt-front: warning Default remote logging level: notice No remote syslog servers configured. Allow receiving of messages from remote hosts: no Number of archived log files to keep: 10 Log rotation size threshold: 5.000% of partition (43 megabytes) Log format: standard Subsecond timestamp field: enabled Subsecond timestamp precision: 1 whole digit; 3 fractional digits Levels at which messages are logged: CLI commands: info Audit messages: notice switch (config) #</pre>
Related Commands	<pre>logging fields logging files rotation logging level logging local logging receive logging <syslog IP address></pre>
Notes	

show log

show log [continues | 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
Role	admin	
Example	<pre>switch (config) # show log matching "Executing Action" Jan 19 10:55:38 arc-switch14 cli28202: [cli.NOTICE]: user admin: Executing command: en Jan 19 11:19:32 arc-switch14 cli28202: [cli.NOTICE]: user admin: Executing command: image install image-SX_PPC_M460EX-ppc-m460ex-20140119-115026.img Jan 19 11:19:32 arc-switch14 mgmtd4064: [mgmtd.NOTICE]: Action ID 326: requested by: user admin (System Administrator) via CLI Jan 19 11:19:32 arc-switch14 mgmtd4064: [mgmtd.NOTICE]: Action ID 326: descr: install system software image Jan 19 11:19:32 arc-switch14 mgmtd4064: [mgmtd.NOTICE]: Action ID 326: param: image file- name: image-SX_PPC_M460EX-ppc-m460ex-20140119-115026.img, version: SX_PPC_M460EX 3.0.0000-dev-master-HA 2014-01-19 11:50:26 ppc Jan 19 11:19:32 arc-switch14 mgmtd4064: [mgmtd.NOTICE]: Action ID 326: param: switch next boot location after install: no switch (config) #</pre>	

Related Commands

logging fields
logging files rotation
logging level
logging local
logging receive
logging <syslog IP address>
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.

4.9 Debugging

➤ *To use the debugging logs feature:*

Step 1. Enable debugging. Run:

```
switch (config) # debug ethernet all
```

Step 2. Display the debug level set. Run:

```
switch (config) # show debug ethernet
```

Step 3. Display the logs. Run:

```
switch (config) # show log debug {match|continue}
```

4.9.1 Commands

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
Role	admin
Example	switch (config) # debug ethernet all switch (config) #
Related Commands	
Notes	

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	
Role	admin	
Example	<pre>switch (config) # debug ethernet dcbx all switch (config) #</pre>	
Related Commands		
Notes		

debug ethernet ip all

debug ethernet ip all

Enables debug traces for all routing modules.

The no form of the command disables debug traces for all routing modules.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.3.4150
Role	admin
Example	switch (config) # debug ethernet ip all switch (config) #
Related Commands	
Notes	

debug ethernet ip arp all

debug ethernet ip arp all
no debug ethernet ip arp all

Enables the trace level for ARP.
The no form of the command disables the trace level for ARP.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.3.4150
Role	admin
Example	switch (config) # debug ethernet ip arp all switch (config) #
Related Commands	
Notes	

debug ethernet ip bgp

debug ethernet ip bgp {all | control-path | dampening | graceful-restart | internal | keep-alive | receive | resources | rtm | transmit | update}
no debug ethernet ip bgp {all | control-path | dampening | graceful-restart | internal | keep-alive | receive | resources | rtm | transmit | update}

Enables the trace level for BGP.

The no form of the command disables tracking a specified level.

Syntax Description	all	Enable track traces
	control-path	Control path dump trace
	dampening	Dampening information
	graceful-restart	Graceful-restart events
	internal	Internal events
	keep-alive	Keep-alive packets exchange
	neighbor	Peer connection/state changes traces
	receive	All received packets
	resources	OS Resource trace
	rtm	Route change notifications
	transmit	All transmitted packets
	update	Update packets exchange
Default	N/A	
Configuration Mode	Config	
History	3.3.4150	
Role	admin	
Example	switch (config) # debug ethernet ip arp all switch (config) #	
Related Commands		
Notes		

debug ethernet ip dhcp-relay

debug ethernet ip dhcp-relay {all | error}
no debug ethernet ip dhcp-relay {all | error}

Configures the trace level for DHCP.
 The no form of the command disables tracking a specified level.

Syntax Description	all	Enables track traces
	error	Error code debug messages
Default	N/A	
Configuration Mode	Config	
History	3.3.4150	
Role	admin	
Example	switch (config) # debug ethernet ip dhcp-relay all switch (config) #	
Related Commands		
Notes		

debug ethernet ip igmp-l3

debug ethernet ip igmp-l3 {all | control-plane | data-path | fail-all | init-shut | management | memory | packet-path | resources}
no debug ethernet ip igmp-l3 {all | control-plane | data-path | fail-all | init-shut | management | memory | packet-path | resources}

Configures the trace level for IGMP.
The no form of the command disables tracking a specified level.

Syntax Description	all	Enable track traces
	control-plane	Control plane traces
	data-path	IP packet dump trace
	fail-all	All failures including Packet Validation Trace
	init-shut	Init and shutdown messages
	management	Management messages
	memory	Memory related messages
	packet-dump	Packet dump messages
	resources	OS resource trace
Default	N/A	
Configuration Mode	Config	
History	3.3.4150	
Role	admin	
Example	switch (config) # debug ethernet ip igmp-l3 all switch (config) #	
Related Commands		
Notes		

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}

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}

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
Default	N/A	
Configuration Mode	Config	
History	3.3.4150	
Role	admin	
Example	<pre>switch (config) # debug ethernet ip igmp-snooping all switch (config) #</pre>	
Related Commands		
Notes		

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	
Role	admin	
Example	switch (config) # debug ethernet ip interface all switch (config) #	
Related Commands		
Notes		

debug ethernet ip ospf

debug ethernet ip ospf {adjacency | all | configuration | ddp-packet | helper | Interface | ism | lrq-packet | lsa_packet | lsu-packet}

Configures the trace level for OSPF.

The no form of the command disables tracking a specified level.

Syntax Description	adjacency	Adjacency formation debug messages
	all	Enable track traces
	configuration	Configuration debug messages
	ddp-packet	DDP packet debug messages
	helper	Helper debug messages
	Interface	Interface debug messages
	ism	Interface State Machine debug messages
	lrq-packet	Link State Request Packet debug messages
	lsa_packet	Link State Acknowledge Packet debug messages
	lsu-packet	Link State Update Packet debug messages
	nsm	Neighbor State Machine debug messages
Default	N/A	
Configuration Mode	Config	
History	3.3.4150	
Role	admin	
Example	switch (config) # debug ethernet ip ospf all switch (config) #	
Related Commands		
Notes		

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	
Role	admin	
Example	switch (config) # debug ethernet lacp all switch (config) #	
Related Commands		
Notes		

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	
Role	admin	
Example	<pre>switch (config) # debug ethernet lldp all switch (config) #</pre>	

Related Commands

Notes

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
Role	admin
Example	switch (config) # debug ethernet port all switch (config) #
Related Commands	
Notes	

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	
Role	admin	
Example	switch (config) # debug ethernet port all switch (config) #	
Related Commands		
Notes		

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
Role	admin
Example	switch (config) # debug ethernet spanning-tree all switch (config) #
Related Commands	
Notes	

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	
Role	admin	
Example	switch (config) # debug ethernet vlan all switch (config) #	
Related Commands		
Notes		

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	
Role	admin	
Example	<pre>switch (config) # show debug ethernet dcbx dcbx protocol : management is ON fail-all is ON control-panel is ON tlv is ON switch (config) #</pre>	
Related Commands		
Notes		

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	
Role	admin	
Example	<pre>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] . . . switch (config) #</pre>	
Related Commands		
Notes		

4.10 Event Notifications

MLNX-OS features a variety of supported events. Events are printed in the system log file, and, optionally, can be sent to the system administrator via email, SNMP trap or directly prompted to the terminal.

4.10.1 Supported Events

Table 28 presents the supported events and maps them to their relevant MIB OID.

For Proxy-ARP related event notifications, please refer to [Section 8.3, “Proxy-ARP Event Notifications,”](#) on page 1556.

Table 28 - Supported Event Notifications and MIB Mapping

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
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
interface-down	An interface's link state has changed to DOWN	RFC1213: linkdown (SNMPv1)	Supported for Ethernet, Infini-Band 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, Infini-Band and management interfaces for 1U and blade systems
internal-bus-error	Internal bus (I ² C) error	Mellanox-EFM-MIB: internalBusError	N/A

Table 28 - Supported Event Notifications and MIB Mapping

Event Name	Event Description	MIB OID	Comments
internal-speed-mismatch	There is a mismatch in the speeds of the internal links between spine and leaf modules	Mellanox-EFM-MIB: internalSpeedMismatch	Relevant for SX65xx and CS75xx switches systems only
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
new_root	Local bridge became a root bridge	Bridge-MIB: newRoot	Supported for Ethernet
paging-high	Paging activity has risen too high	N/A	Not supported
power-redundancy-mismatch	Power redundancy mismatch	Mellanox-EFM-MIB: powerRedundancyMismatch	Supported only for director switch systems
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
snmp-authtrap	An SNMPv3 request has failed authentication	Not implemented	N/A
topology_change	Topology change triggered by a local bridge	Bridge-MIB: topologyChange	Supported for Ethernet
unexpected-shutdown	Unexpected system shutdown	Mellanox-EFM-MIB: unexpectedShutdown	N/A
To send, use the CLI command <code>snmp-server notify send-test</code>	Send a testing event	testTrap	N/A
N/A	Reset occurred due to over-heating of ASIC	Mellanox-EFM-MIB: asicOverTempReset	Not supported
temperature-too-high	Temperature is too high	Mellanox-EFM-MIB: asicOverTemp	N/A

4.10.2 SNMP Trap Notifications

To set SNMP notification see [Section 4.18.1.6, “Configuring an SNMP Notification,”](#) on page 551.

4.10.3 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 2013: Interface IB1/17 changed state to DOWN  
Wed Jul 10 11:30:43 2013: Interface IB1/18 changed state to DOWN  
switch (config) #
```

4.10.4 Email Notifications

➤ *To configure MLNX-OS to send you emails for all configured events and failures:*

Step 1. Enter to Config mode. Run:

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

Step 2. 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>
```

Step 3. Add your email address for notifications. Run:

```
switch (config) # email notify recipient <email address>
```

Step 4. Configure the system to send notifications for a specific event. Run:

```
switch (config) # email notify event <event name>
```

Step 5. 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  
  .  
  .  
  .  
switch (config) #
```

Step 6. Have the system send you a test email. Run:

```
switch # 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.

4.10.5 Commands

4.10.5.1 Email Notification

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
Role	admin
Example	switch (config) # email autosupport enable
Related Commands	N/A
Notes	

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
Role	admin
Example	switch (config) # email autosupport event process-crash
Related Commands	N/A
Notes	

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 plain-text if this fails.
Default	tls-none	
Configuration Mode	Config	
History	3.2.3000	
Role	admin	
Example	switch (config) # email autosupport ssl mode tls	
Related Commands	N/A	
Notes		

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
Role	admin
Example	switch (config) # email autosupport ssl cert-verify
Related Commands	N/A
Notes	

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	
Role	admin	
Example	switch (config) # email autosupport ssl ca-list default_ca_list	
Related Commands	N/A	
Notes		

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	
Role	admin	
Example	switch (config) # email dead-letter enable switch (config) #	
Related Commands	show email	
Notes		

email domain

email domain <hostname or IP address>
no email domain

Sets the domain name from which the emails will appear to come from (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 IP address.
Default	No email domain
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # email domain mellanox switch (config) # show email Mail hub: 10.0.8.11 Mail hub port: 125 Domain: mellanox Return address: do-not-reply Include hostname in return address: yes ... switch (config) #</pre>
Related Commands	show emails
Notes	

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
Role	admin
Example	<pre>switch (config) # email mailhub 10.0.8.11 switch (config) # show email Mail hub: 10.0.8.11 Mail hub port: 25 Domain: (not specified) Return address: do-not-reply Include hostname in return address: yes ... switch (config) #</pre>
Related Commands	show email [events]
Notes	

email mailhub-port

email mailhub-port <hostname or IP address>
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	hostname or IP address.
Default	25	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # email mailhub-port 125 switch (config) # show email Mail hub: 10.0.8.11 Mail hub port: 125 Domain: (system domain name) Return address: do-not-reply Include hostname in return address: yes ... switch (config) #</pre>	
Related Commands	show email	
Notes		

email notify event

email notify event <event name>
no email notify event <event name>

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 name	Example event names would include “process-crash” and “cpu-util-high”.
Default	No events are enabled	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # email notify event process-crash 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: liveness-failure: A process in the system was detected as hung process-exit: A process in the system unexpectedly exited cpu-util-ok: CPU utilization has fallen back to normal levels cpu-util-high: CPU utilization has risen too high disk-io-ok: Disk I/O per second has fallen back to acceptable levels ... temperature-too-high: Temperature has risen too high All events for which autosupport emails will be sent: process-crash: A process in the system has crashed liveness-failure: A process in the system was detected as hung switch (config) # switch (config) #</pre>	
Related Commands	show email	
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.	

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	No recipients are added	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # email notify recipient user2@autosupport.mellanox.com switch (config) # show email Mail hub: Mail hub port: 25 Domain: (not specified) Return address: user1 Include hostname in return address: no Dead letter settings: Save dead.letter files: yes Dead letter max age: (none) Email notification recipients: user2@autosupport.mellanox.com (all events, in detail) Autosupport emails Enabled: no Recipient: autosupport@autosupport.mellanox.com Mail hub: autosupport.mellanox.com switch (config) #</pre>	
Related Commands	show email	
Notes		

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	do-not-reply
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre> switch (config) # email return-addr user1 switch (config) # show email Mail hub: Mail hub port: 25 Domain: (not specified) Return address: user1 Include hostname in return address: yes ... switch (config) # </pre>
Related Commands	show email
Notes	

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
Role	admin
Example	<pre> switch (config) # no email return-host switch (config) # show email Mail hub: Mail hub port: 25 Domain: (system domain name) Return address: my-address Include hostname in return address: no Current reply address: host@localdomain Dead letter settings: Save dead.letter files: yes Dead letter max age: 5 days No recipients configured. Autosupport emails Enabled: no Recipient: autosupport@autosupport.mellanox.com Mail hub: autosupport.mellanox.com switch (config) # </pre>
Related Commands	show email
Notes	This only takes effect if the return address does not contain an “@” character.

email send-test

email send-test

Sends test-email to all configured event and failure recipients.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # email send-test
Related Commands	show email [events]
Notes	

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	
Role	admin	
Example	switch (config) # email ssl mode tls-none	
Related Commands	N/A	
Notes		

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
Role	admin
Example	switch (config) # email ssl cert-verify
Related Commands	N/A
Notes	This command has no impact unless TLS is used.

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	
Role	admin	
Example	switch (config) # email ssl ca-list none	
Related Commands	N/A	
Notes	This command has no impact unless TLS is used, and certificate verification is enabled.	

show email

show email [events]

Shows email configuration or events for which email should be sent upon.

Syntax Description	events	show event list
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # show email Mail hub: Mail hub port: 25 Domain: (system domain name) Return address: my-address Include hostname in return address: no Current reply address: host@localdomain Dead letter settings: Save dead.letter files: yes Dead letter max age: 5 days No recipients configured. Autosupport emails Enabled: no Recipient: autosupport@autosupport.mellanox.com Mail hub: autosupport.mellanox.com switch (config) #</pre>	
Related Commands	show email	
Notes		

4.11 Telemetry

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.

MLNX-OS provides following mechanisms to perform those 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 – recording of most active flows which cause an excessive usage of the port buffers. 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.

4.11.1 Commands

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	Disabled.
Configuration Mode	Config
History	3.6.3004
Role	admin
Example	switch (config) # protocol telemetry switch (config) # no protocol telemetry

Related Commands

Notes

telemetry shutdown

telemetry shutdown
no telemetry shutdown

Disables the telemetry protocol, 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
Role	admin
Example	switch (config) # telemetry shutdown switch (config) # no telemetry shutdown
Related Commands	
Notes	

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 msec - 1 min
Default	1000 msec.	
Configuration Mode	Config	
History	3.6.3004	
Role	admin	
Example	switch (config) # telemetry sampling log 1000 switch (config) # no telemetry sampling log	
Related Commands		
Notes		

interface ethernet telemetry sampling tc mcast

```
interface ethernet <slot>/<port>[/<subport>] telemetry sampling tc <tc_id>
mcast
no interface ethernet <slot>/<port>[/<subport>] telemetry sampling tc <tc_id>
mcast
```

Enables multicast sampling (histogram fetching) on a tc for a specific Ethernet interface.

The no form of the command disables multicast sampling on a tc for a specific Ethernet interface.

Syntax Description	slot/port	Ethernet port number
	subport	Ethernet subport number to be used in case of split port
	tc_id	Input range: 0-7
Default	N/A	
Configuration Mode	Config-Instance-ID	
History	3.6.3004 1/2/4	
Role	admin	
Example	switch (config) # interface ethernet 1/2 telemetry sampling tc 3 mcast switch (config) # no interface ethernet 1/2 telemetry sampling tc 3 mcast	
Related Commands		
Notes		

interface ethernet telemetry sampling tc ucast

```
interface ethernet <slot>/<port>[/<subport>] telemetry sampling tc <tc_id>
ucast
no interface ethernet <slot>/<port>[/<subport>] telemetry sampling tc <tc_id>
ucast
```

Enables unicast sampling (histogram fetching) on a tc for a specific Ethernet interface.

The no form of the command disables unicast sampling on a tc for a specific Ethernet interface.

Syntax Description	slot/port	Ethernet port number
	subport	Ethernet subport number to be used in case of split port
	tc_id	Input range: 0-7
Default	N/A	
Configuration Mode	Config	
History	3.6.3004	
Role	admin	
Example	switch (config) # interface ethernet 1/2 telemetry sampling tc 3 ucast switch (config) # no interface ethernet 1/2 telemetry sampling tc 3 ucast	
Related Commands		
Notes		

interface ib telemetry sampling

interface ib <slot>/<port> telemetry sampling
no interface ib <slot>/<port> telemetry sampling

Enables sampling (histogram fetching) for a specific InfiniBand interface.
 The no form of the command disables sampling (histogram fetching).

Syntax Description	slot/port	Infiniband port number
Default	N/A	
Configuration Mode	Config-Instance	
History	3.6.3004	
Role	admin	
Example	switch (config) # interface ib 1/1 telemetry sampling switch (config) # no interface ib 1/1 telemetry sampling	
Related Commands		
Notes		

debug generate teledump

fae telemetry sampling binsize min <bin_size> resolution <bin_resolution>
no fae telemetry sampling binsize min <bin_size> resolution <bin_resolution>

Hidden command which dumps telemetry cache in dump file.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.6.3004
Role	admin
Example	switch (config) # debug generate teledump dump action successful. dump file is at: /var/log/teledump
Related Commands	
Notes	

show telemetry sampling interface ethernet tc ucast

show telemetry sampling interface ethernet <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 in case of split port																																												
	tc_id	Input range: 0-7																																												
Default	N/A																																													
Configuration Mode	Any Command Mode																																													
History	3.6.3004																																													
Role	admin																																													
Example	<pre>switch (config) # show telemetry sampling interface ethernet 1/2 tc 6 ucast</pre> Telemetry histogram: Eth1/2 traffic-class 6 - ucast TimeBin sizes (nsec buffer was occupied in bytes range) <table><tr><td>01/13/17</td><td>2976<</td><td>27552</td><td>52128</td><td>76704</td><td>101280</td><td>125856</td><td>150432</td><td>175008</td><td>199584</td><td>199584></td></tr><tr><td>08:18:09.67745</td><td>1000000000</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td></tr><tr><td>08:18:10.67850</td><td>1000000000</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td></tr><tr><td>08:18:11.67953</td><td>1000000000</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td></tr></table>		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
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																																														

show telemetry sampling interface ethernet tc mcast

show telemetry sampling interface ethernet <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 in case of split port																																																																													
	tc_id	Input range: 0-7																																																																													
Default	N/A																																																																														
Configuration Mode	Any Command Mode																																																																														
History	3.6.3004																																																																														
Role	admin																																																																														
Example	<pre>switch (config) # show telemetry sampling interface ethernet 1/2 tc 3 mcast</pre> Telemetry histogram: Eth1/2 traffic-class 3 - mcast <table><thead><tr><th>Time</th><th colspan="10">Bin sizes (nsec buffer was occupied in bytes range)</th></tr></thead><tbody><tr><td>01/16/17 2976<</td><td>27552</td><td>52128</td><td>76704</td><td>101280</td><td>125856</td><td>150432</td><td>175008</td><td>199584</td><td>199584></td><td></td></tr><tr><td>04:09:07.79936 1000000000</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td></td></tr><tr><td>04:09:08.80096 1000000000</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td></td></tr><tr><td>04:09:09.80355 1000000000</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td></td></tr><tr><td>04:09:10.80518 1000000000</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td></td></tr><tr><td>04:09:11.80682 1000000000</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td></td></tr></tbody></table> 		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	
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																																																																															

show telemetry sampling interface ethernet tc ucast last

show telemetry sampling interface ethernet <slot>/<port>[/<subport>] tc <tc_id> ucast last <num_of_entries>

Displays last num of fetched unicast 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 in case of split port																																																							
	tc_id	Input range: 0-7																																																							
	num_of_entries	Input range: 0-1000																																																							
Default	N/A																																																								
Configuration Mode	Any Command Mode																																																								
History	3.6.3004																																																								
Role	admin																																																								
Example	<pre>switch (config) # show telemetry sampling interface ethernet 1/2 tc 3 ucast last 3</pre> Telemetry histogram: Eth1/2 traffic-class 3 - ucast <table><thead><tr><th>Time</th><th colspan="10">Bin sizes (nsec buffer was occupied in bytes range)</th></tr></thead><tbody><tr><td>01/16/17</td><td>2976<</td><td>27552</td><td>52128</td><td>76704</td><td>101280</td><td>125856</td><td>150432</td><td>175008</td><td>199584</td><td>199584></td></tr><tr><td>04:28:39.81351</td><td>1000000000</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td></tr><tr><td>04:28:40.81512</td><td>1000000000</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td></tr><tr><td>04:28:41.81708</td><td>1000000000</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td></tr></tbody></table>		Time	Bin sizes (nsec buffer was occupied in bytes range)										01/16/17	2976<	27552	52128	76704	101280	125856	150432	175008	199584	199584>	04:28:39.81351	1000000000	0	0	0	0	0	0	0	0	0	04:28:40.81512	1000000000	0	0	0	0	0	0	0	0	0	04:28:41.81708	1000000000	0	0	0	0	0	0	0	0	0
Time	Bin sizes (nsec buffer was occupied in bytes range)																																																								
01/16/17	2976<	27552	52128	76704	101280	125856	150432	175008	199584	199584>																																															
04:28:39.81351	1000000000	0	0	0	0	0	0	0	0	0																																															
04:28:40.81512	1000000000	0	0	0	0	0	0	0	0	0																																															
04:28:41.81708	1000000000	0	0	0	0	0	0	0	0	0																																															
Related Commands																																																									
Notes	In case requested entries are more than what the DB contains it will print the amount in the table.																																																								

show telemetry sampling interface ethernet tc mcast last

show telemetry sampling interface ethernet <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 in case of split port																																																																													
	tc_id	Input range: 0-7																																																																													
	num_of_entries	Input range: 0-1000																																																																													
Default	N/A																																																																														
Configuration Mode	Any Command Mode																																																																														
History	3.6.3004																																																																														
Role	admin																																																																														
Example	<pre>switch (config) # show telemetry sampling interface ethernet 1/2 tc 3 mcast last 4</pre> Telemetry histogram: Eth1/2 traffic-class 3 - mcast <table><thead><tr><th>Time</th><th colspan="10">Bin sizes (nsec buffer was occupied in bytes range)</th></tr><tr><td colspan="11">.....</td></tr></thead><tbody><tr><td>01/16/17</td><td>2976<</td><td>27552</td><td>52128</td><td>76704</td><td>101280</td><td>125856</td><td>150432</td><td>175008</td><td>199584</td><td>199584></td></tr><tr><td>04:23:38.28864</td><td>1000000000</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td></tr><tr><td>04:23:39.28977</td><td>1000000000</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td></tr><tr><td>04:23:40.29111</td><td>1000000000</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td></tr><tr><td>04:23:41.29259</td><td>1000000000</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td></tr></tbody></table>		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
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	In case requested entries are more than what the DB contains it will print the amount in the table.																																																																														

show telemetry sampling interface ib

show telemetry sampling interface ib <slot>/<port>

Displays fetched histogram details for tc_id of an Ethernet interface.

Syntax Description	slot/port	Infiniband port number
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.6.3004	
Role	admin	
Example	<pre>switch (config) # show telemetry sampling interface ethernet 1/2 tc 3 Telemetry histogram: Eth1/2 traffic-class 3 - anycast Time Bin sizes (nsec buffer was occupied in bytes range) 01/13/17 2976< 27552 52128 76704 101280 125856 150432 175008 199584 199584> 07:50:13.66895 976563 0 0 0 0 0 0 0 0 07:50:14.66932 976563 0 0 0 0 0 0 0 0 07:50:15.66990 976563 0 0 0 0 0 0 0 0</pre>	
Related Commands		
Notes	In case requested entries are more than what the DB contains it will print the amount in the table.	

show telemetry sampling interface ib last

show telemetry sampling interface ib <slot>/<port> last <num_of_entries>

Displays fetched unicast histogram details for tc_id of an Ethernet interface.

Syntax Description	slot/port	Infiniband port number
	num_of_entries	
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.6.3004	
Role	admin	
Example	<pre>switch (config) # show telemetry sampling interface ethernet 1/2 tc 6 ucast</pre> <pre>-----</pre> <pre>Telemetry histogram: Eth1/2 traffic-class 6 - ucast</pre> <pre>Time Bin sizes (nsec buffer was occupied in bytes range)</pre> <pre>-----</pre> <pre>01/13/17 2976< 27552 52128 76704 101280 125856 150432 175008 199584 199584></pre> <pre>08:18:09.67745 976563 0 0 0 0 0 0 0 0</pre> <pre>08:18:10.67850 976563 0 0 0 0 0 0 0 0</pre> <pre>08:18:11.67953 976563 0 0 0 0 0 0 0 0</pre>	
Related Commands		
Notes	In case requested entries are more than what the DB contains it will print the amount in the table.	

stats export csv telemetry

stats export csv telemetry <slot>[/<port>[/<subport>] [filename *] [after * *]
[before * *]

Exports histograms collected by stats to a csv file.

Syntax Description	slot/port	Ethernet Infiniband port number
	subport	Ethernet Infiniband subport number to be used in case of split port
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.6.3004	
Role	admin	
Example	switch (config) # stats export csv telemetry 1/1 Generated report file: telemetry-20170119-102715.csv	
Related Commands		
Notes		

file stats telemetry upload

file stats telemetry upload <filename> <upload URL>

Uploads file created by stats export command to user directory.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.6.3004
Role	admin
Example	<pre>switch (config) # file stats telemetry upload telemetry-20170119-102715.csv scp://username:password@server//directory</pre>
Related Commands	
Notes	

show files stats telemetry

show files stats telemetry [filename]

Displays all files created by command stats export csv telemetry unless a filename is given.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.6.3004
Role	admin
Example	switch (config) # show files stats telemetry telemetry-20170119-102715.csv
Related Commands	
Notes	

4.12 mDNS

Multicast DNS (mDNS) protocol is used by the SM HA to deliver control information between the InfiniBand nodes via the management interface. To block sending mDNS traffic from the management interface run the command `no ha dns enable`.

4.12.1 Commands

ha dns enable

ha dns enable
no ha dns enable

Allows mDNS traffic.
 The no form of the command blocks mDNS traffic from being sent from mgmt0.

Syntax Description	N/A
Default	Enabled.
Configuration Mode	Config
History	3.3.4000
Role	admin
Example	switch (config) # no ha dns enable switch (config) #
Related Commands	
Notes	

4.13 User Management and Security

4.13.1 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.

Table 29 - User Roles (Accounts) and Default Passwords

User Role	Default Password
admin	admin
monitor	monitor
xmladmin	xmladmin
xmluser	xmluser

To remove passwords from the XML users, run the command `username <username> nopassword`.

4.13.2 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 MLNX-OS switch. The MLNX-OS switch supports Remote Access Dial-In User Service (RADIUS) or Terminal Access Controller Access Control device Plus (TACACS+) protocols.

- **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.

4.13.2.1 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 etc.) that has been previously used to authenticate a user is modified, that user gets immediately logged out of the switch and asked to re-authenticate.

4.13.2.2 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

4.13.2.3 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

4.13.2.4 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 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. If the user-account name is John, the following is an example DN:

```
uid=John,ou=people,dc=domain,dc=com
```

4.13.3 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:

Step 1. Disable SNMPv1 and v2. Run:

```
switch (config) # no snmp-server enable communities
```

Step 2. Only allow SNMPv3 users with sha and aes-128. Run:

```
switch (config) # snmp-server user <username> v3 auth sha <password1> priv aes-128  
<password2>
```

Step 3. Only allow SNMPv3 traps with sha and aes-128. Run:

```
switch (config) # snmp-server host <ip-address> informs version 3 user <username> auth  
sha <password1> priv aes-128 <password2>
```

Step 4. Only allow SSHv2. Run:

```
switch (config) # ssh server min-version 2
```

Step 5. Enable SSH server strict security mode. Run:

```
switch (config) # ssh server security strict
```

Step 6. Disable HTTP access. Run:

```
switch (config) # no web http enable
```

Step 7. Enable HTTPS strict cyphers. Run:

```
switch (config) # web https ssl ciphers TLS1.2
```

Step 8. Disable router BGP neighbor password configuration. Run:

```
switch (config) # no router bgp <as-number> neighbor <ip-address> password
```

Step 9. Disable router BGP peer group password configuration. Run:

```
switch (config) # no router bgp <as-number> peer-group <peer-group-name> password
```

Step 10. Disable BGP password configuration. Run:

```
switch (config) # no neighbor <ip-address> password
```

Step 11. Disable MD5 password hashing on for users. Run:

```
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.



Secure mode is not supported on director switch systems.

➤ **To activate secure mode:**

```
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:**

```
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:**

```
switch (config)# show system secure-mode
```

```
Secure mode configured: yes
Secure mode enabled: yes
switch (config) #
```

4.13.4 Commands

4.13.4.1 User Accounts

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.
capability <cap>	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
disable [login password]	- Disable - disable this account - Disable login - disable all logins to this account - Disable password - disable login to this account using a local password
disconnect	Logs out the specified user from the system
name	Full name of the user
nopassword	The next login of the user will not require password.
0 7	0: specifies a login password in cleartext 7: specifies a login password in encrypted text
password	Specifies a password for the user in string form. If [0 7] was not specified then the password is in cleartext.
Default	The following usernames are available by default: - admin - monitor - xmladmin - xmluser

Configuration Mode	Config
History	3.1.0000 3.4.0000 Updated Example 3.4.1100 Updated Example 3.6.2002 Added “disconnect” parameter
Role	admin
Example	<pre>switch (config) # username monitor full-name smith switch (config) # show usernames USERNAME FULL NAME CAPABILITY ACCOUNT STATUS USERID System Administrator admin Password set admin System Administrator admin Password set monitor smith monitor Password set (SHA512) xmladmin XML Admin User admin Password set (SHA512) xmluser XML Monitor User monitor Password set (SHA512) switch (config) #</pre>
Related Commands	show usernames show users
Notes	• To enable a user account, just set a password on it (or use the command <code>username <user> nopassword</code> 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 <code>show configuration</code>, since the cleartext password cannot be recovered after it is set

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
Role	admin
Example	<pre>switch (config) # show usernames USERNAME FULL NAME CAPABILITY ACCOUNT STATUS USERID admin System Administrator admin Password set monitor smith monitor Password set (SHA512) xmladmin XML Admin User admin No password required xmluser XML Monitor User monitor No password required switch (config) #</pre>
Related Commands	<pre>username show users</pre>
Notes	

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	
Role	admin	
Example	<pre>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) #show 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 switch (config) #</pre>	
Related Commands	username show usernames	
Notes		

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
Role	admin
Example	switch (config) # show whoami Current user: admin Capabilities: admin switch (config) #
Related Commands	username show usernames show users
Notes	

4.13.4.2 AAA Methods

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 First version 3.2.3000 Removed 'time' parameter from the command.
Role	admin
Example	<pre>switch (config) # aaa accounting changes default stop-only tacacs+ switch (config) # show aaa AAA authorization: Default User: admin Map Order: local-only Authentication method(s): local radius tacacs+ ldap Accounting method(s): tacacs+ switch (config) #</pre>
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 “logging level audit mgmt” or “configuration audit” commands • 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 data-base changes, system actions are logged when the action is started, not when the action has completed

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
Role	admin
Example	<pre>switch (config) # aaa authentication login default local radius tacacs+ ldap switch (config) # show aaa AAA authorization: Default User: admin Map Order: local-only Authentication method(s): local radius tacacs+ ldap Accounting method(s): tacacs+ switch (config) #</pre>
Related Commands	show aaa
Notes	The order in which the methods are specified is the order in which the authentication is attempted. It is required that “local” is one of the methods selected. It is recommended that “local” be listed first to avoid potential problems logging in to local accounts in the face of network or remote server issues.

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	
Role	admin	
Example	switch (config) # aaa authentication attempts fail-delay 1	
Related Commands	N/A	
Notes		

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	3.2.3000	
Role	admin	
Example	switch (config) # aaa authentication attempts track enable	
Related Commands	N/A	
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.	

aaa authentication attempts logout

```
aaa authentication attempts logout {enable | lock-time | max-fail | unlock-time}  
no aaa authentication attempts logout {enable | lock-time | max-fail | unlock-time}
```

Configures logout of accounts based on failed authentication attempts.
The no form of the command clears configuration for logout 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 retro-active 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_tally-byname command-line utility to unlock your

Default	N/A
Configuration Mode	Config
History	3.2.3000
Role	admin
Example	switch (config) # aaa authentication attempts logout enable
Related Commands	N/A
Notes	

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 the 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	
Role	admin	
Example	switch (config) # aaa authentication attempts class-override admin no-lockout	

Related Commands	N/A

Notes

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 specified 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	
Role	admin	
Example	switch (config) # aaa authentication attempts reset user admin all	
Related Commands	N/A	
Notes		

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 specified 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	
Role	admin	
Example	switch (config) # aaa authentication attempts reset user admin no-clear-history	
Related Commands	N/A	
Notes		

aaa authorization

aaa authorization map [default-user <username> | order <policy>]
no aaa authorization map [default-user | order]

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+) 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+ 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: - 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. - local-only – maps all remote users to the user specified by the “aaa authorization map default-user <user name>” command. Any vendor attributes received by an authentication server are ignored.
Default	Default user - admin Map order - remote-first	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	

Example

```
switch (config) # aaa authorization map default-user admin
switch (config) # show aaa
AAA authorization:
  Default User: admin
  Map Order: remote-first
Authentication method(s):
  local
Accounting method(s):
  tacacs+
switch (config) #
```

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.
 - If AAA authorization order policy is configured to remote-only, then when upgrading to 3.4.3000 or later from an older MLNX-OS version, this policy is changed to remote-first.
 - The user must be careful when setting AAA authorization to “remote-only” because if the remote server happens to be configured incorrectly, then the user may lock themselves out.
-

show aaa

show aaa

Displays the AAA configuration.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.1.0000
Role	admin
Example	<pre>switch (config) # show aaa AAA authorization: Default User: admin Map Order: remote-first Authentication method(s): local Accounting method(s): tacacs+ switch (config) #</pre>
Related Commands	<pre>aaa accounting aaa authentication aaa authorization show aaa show usernames username</pre>
Notes	

show aaa authentication attempts

show aaa authentication attempts [configured | status user <username>]]

Shows 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
Role	admin	
Example	<pre>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 2012/08/20 14:29:19 ttyS0 (*) Hashed for security reasons switch-627d3c [standalone: master] (config) # switch (config) #</pre>	
Related Commands	N/A	
Notes		

4.13.4.3 RADIUS

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	
Role	admin	
Example	<pre>switch (config) #radius-server retransmit 3 switch (config) # show radius RADIUS defaults: Key: 3333 Timeout: 3 Retransmit: 1 No RADIUS servers configured. switch (config) #</pre>	
Related Commands	<pre>aaa authorization radius-server host show radius</pre>	
Notes	Each RADIUS server can override those global parameters using the command “radius-server host”.	

radius-server host

**radius-server host <IP address> [enable | auth-port <port> | key <secret> | prompt-key | retransmit <retries> | timeout <seconds>]
no radius-server host <IP address> [auth-port | enable]**

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)
Default	3 seconds, 1 retry Default UDP port is 1812
Configuration Mode	Config
History	3.1.0000
Role	admin

Example

```
switch (config) # radius-server host 40.40.40.40
switch (config) # show radius
RADIUS defaults:
  Key:          3333
  Timeout:      3
  Retransmit:   1
RADIUS servers:
  40.40.40.40:1812
    Enabled:    yes
    Key:        3333 (default)
    Timeout:    3 (default)
    Retransmit: 1 (default)
switch (config) #
```

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 “radius-server” command.
-

show radius

show radius

Displays RADIUS configurations.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.1.0000
Role	admin
Example	<pre> switch (config) # show radius RADIUS defaults: Key: 3333 Timeout: 3 Retransmit: 1 RADIUS servers: 40.40.40.40:1812 Enabled: yes Key: 3333 (default) Timeout: 3 (default) Retransmit: 1 (default) switch (config) # </pre>
Related Commands	<pre> aaa authorization radius-server radius-server host </pre>
Notes	

4.13.4.4 TACACS+

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 (1-60).
Default	3 seconds, 1 retry	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) #tacacs-server retransmit 3 switch (config) # show tacacs TACACS+ defaults: Key: 3333 Timeout: 3 Retransmit: 1 No TACACS+ servers configured. switch (config) #</pre>	
Related Commands	<pre>aaa authorization show radius show tacacs tacacs-server host</pre>	
Notes	Each TACACS+ server can override those global parameters using the command “tacacs-server host”.	

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 (1-60)
Default	3 seconds, 1 retry Default TCP port is 49 Default auth-type is PAP
Configuration Mode	Config

History	3.1.0000
Role	admin
Example	<pre> switch (config) # tacacs-server host 40.40.40.40 switch (config) # show tacacs TACACS+ defaults: Key: 3333 Timeout: 3 Retransmit: 1 TACACS+ servers: 40.40.40.40:49 Enabled: yes Auth-type PAP Key: 3333 (default) Timeout: 3 (default) Retransmit: 1 (default) switch (config) # </pre>
Related Commands	<pre> aaa authorization show tacacs tacacs-server </pre>
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 “tacacs-server” command.

show tacacs

show tacacs

Displays TACACS+ configurations.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.1.0000
Role	admin
Example	<pre> switch (config) # show tacacs TACACS+ defaults: Key: 3333 Timeout: 3 Retransmit: 1 TACACS+ servers: 40.40.40.40:49 Enabled: yes Auth-type PAP Key: 3333 (default) Timeout: 3 (default) Retransmit: 1 (default) switch (config) # </pre>
Related Commands	<pre> aaa authorization tacacs-server tacacs-server host </pre>
Notes	

4.13.4.5 LDAP

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. when: 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
Role	admin	

Example

```
switch (config) # ldap base-dn ou=department,dc=example,dc=com
switch (config) # show ldap
User base DN      : ou=department,dc=example,dc=com
User search scope : subtree
Login attribute   : sAMAccountName
Bind DN           :
Bind password     :
Group base DN     :
Group attribute   : member
LDAP version      : 3
Referrals         : yes
Server port       : 389
Search Timeout    : 5
Bind Timeout      : 5
SSL mode          : none
Server SSL port   : 636 (not active)
SSL ciphers       : TLS1.2 (not active)
SSL cert verify   : yes
SSL ca-list       : default-ca-list

LDAP servers:
  1: 10.10.10.10
  2: 10.10.10.12
switch (config) #
```

Related Commands

show ldap

Notes

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
Role	admin	
Example	<pre>switch (config) # ldap bind-dn my-dn switch (config) # ldap bind-password my-password switch (config) # show ldap User base DN : ou=department,dc=example,dc=com User search scope : subtree Login attribute : sAMAccountName Bind DN : my-dn Bind password : my-password Group base DN : Group attribute : member LDAP version : 3 Referrals : yes Server port : 389 Search Timeout : 5 Bind Timeout : 5 SSL mode : none Server SSL port : 636 (not active) SSL ciphers : TLS1.2 (not active) SSL cert verify : yes SSL ca-list : default-ca-list LDAP servers: 1: 10.10.10.10 2: 10.10.10.12 switch (config) #</pre>	
Related Commands	show ldap	
Notes	For anonymous login, bind-dn and bind-password should be empty strings "".	

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
Role	admin	
Example	<pre> switch (config) # ldap group-attribute member switch (config) # ldap group-dn my-group-dn switch (config) # show ldap User base DN : ou=department,dc=example,dc=com User search scope : subtree Login attribute : SAMAccountName Bind DN : my-dn Bind password : my-password Group base DN : my-group-dn Group attribute : member LDAP version : 3 Referrals : yes Server port : 389 Search Timeout : 5 Bind Timeout : 5 SSL mode : none Server SSL port : 636 (not active) SSL ciphers : TLS1.2 (not active) SSL cert verify : yes SSL ca-list : default-ca-list LDAP servers: 1: 10.10.10.10 2: 10.10.10.12 switch (config) # </pre>	

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.
-

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
Role	admin	
Example	<pre> switch (config) # ldap host 10.10.10.10 switch (config) # show ldap User base DN : ou=department,dc=example,dc=com User search scope : subtree Login attribute : SAMAccountName Bind DN : my-dn Bind password : my-password Group base DN : my-group-dn Group attribute : member LDAP version : 3 Referrals : yes Server port : 389 Search Timeout : 5 Bind Timeout : 5 SSL mode : none Server SSL port : 636 (not active) SSL ciphers : TLS1.2 (not active) SSL cert verify : yes SSL ca-list : default-ca-list LDAP servers: 1: 10.10.10.10 2: 10.10.10.12 switch (config) # </pre>	

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
-

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
Role	admin	
Example	<pre> switch (config) # ldap login-attribute uid switch (config) # show ldap User base DN : ou=department,dc=example,dc=com User search scope : subtree Login attribute : uid Bind DN : my-dn Bind password : my-password Group base DN : my-group-dn Group attribute : member LDAP version : 3 Referrals : yes Server port : 389 Search Timeout : 5 Bind Timeout : 5 SSL mode : none Server SSL port : 636 (not active) SSL ciphers : TLS1.2 (not active) SSL cert verify : yes SSL ca-list : default-ca-list LDAP servers: 1: 10.10.10.10 2: 10.10.10.12 switch (config) # </pre>	
Related Commands	<pre> show aaa show ldap </pre>	
Notes		

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
Role	admin	
Example	<pre>switch (config) # ldap port 1111 switch (config) # show ldap User base DN : ou=department,dc=example,dc=com User search scope : subtree Login attribute : uid Bind DN : my-dn Bind password : my-password Group base DN : my-group-dn Group attribute : member LDAP version : 3 Referrals : yes Server port : 1111 Search Timeout : 5 Bind Timeout : 5 SSL mode : none Server SSL port : 636 (not active) SSL ciphers : TLS1.2 (not active) SSL cert verify : yes SSL ca-list : default-ca-list LDAP servers: 1: 10.10.10.10 2: 10.10.10.12 switch (config) #</pre>	
Related Commands	<pre>show aaa show ldap</pre>	
Notes		

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
Role	admin
Example	<pre> switch (config) # no ldap referrals switch (config) # show ldap User base DN : ou=department,dc=example,dc=com User search scope : subtree Login attribute : uid Bind DN : my-dn Bind password : my-password Group base DN : my-group-dn Group attribute : member LDAP version : 3 Referrals : no Server port : 1111 Search Timeout : 5 Bind Timeout : 5 SSL mode : none Server SSL port : 636 (not active) SSL ciphers : TLS1.2 (not active) SSL cert verify : yes SSL ca-list : default-ca-list LDAP servers: 1: 10.10.10.10 2: 10.10.10.12 switch (config) # </pre>
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.

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
Role	admin	
Example	<pre> switch (config) # ldap scope subtree switch (config) # show ldap User base DN : ou=department,dc=example,dc=com User search scope : subtree Login attribute : uid Bind DN : my-dn Bind password : my-password Group base DN : my-group-dn Group attribute : member LDAP version : 3 Referrals : no Server port : 1111 Search Timeout : 5 Bind Timeout : 5 SSL mode : none Server SSL port : 636 (not active) SSL ciphers : TLS1.2 (not active) SSL cert verify : yes SSL ca-list : default-ca-list LDAP servers: 1: 10.10.10.10 2: 10.10.10.12 switch (config) # </pre>	
Related Commands	<pre> show aaa show ldap </pre>	
Notes		

ldap ssl

```
ldap ssl {ca-list <options> | cert-verify | ciphers {all | TLS1.2} | mode <mode> |  
port <port-number>}  
no ldap ssl {cert-verify | ciphers | 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.
	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
	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	First version
	3.2.3000	Added ca-list argument.
	3.4.0000	Added “ssl ciphers” parameter Updated Example
Role	admin	
Example	<pre> switch (config) # ldap ssl mode ssl switch (config) # show ldap User base DN : ou=department,dc=example,dc=com User search scope : subtree Login attribute : uid Bind DN : my-dn Bind password : my-password Group base DN : my-group-dn Group attribute : member LDAP version : 3 Referrals : no Server port : 1111 Search Timeout : 5 Bind Timeout : 5 SSL mode : ssl Server SSL port : 636 (not active) SSL ciphers : TLS1.2 (not active) SSL cert verify : yes SSL ca-list : default-ca-list LDAP servers: 1: 10.10.10.10 2: 10.10.10.12 switch (config) # </pre>	
Related Commands	<pre> show aaa show ldap </pre>	
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. In case the mode is TLS, the LDAP port number will be used.	

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	Range: 1-60 seconds.
Default	5 seconds	
Configuration Mode	Config	
History	3.1.0000	
	3.4.0000	Updated Example
Role	admin	
Example	<pre> switch (config) # ldap timeout-bind 10 switch (config) # show ldap User base DN : ou=department,dc=example,dc=com User search scope : subtree Login attribute : uid Bind DN : my-dn Bind password : my-password Group base DN : my-group-dn Group attribute : member LDAP version : 3 Referrals : no Server port : 1111 Search Timeout : 5 Bind Timeout : 10 SSL mode : none Server SSL port : 636 (not active) SSL ciphers : TLS1.2 (not active) SSL cert verify : yes SSL ca-list : default-ca-list LDAP servers: 1: 10.10.10.10 2: 10.10.10.12 switch (config) # </pre>	

Related Commands

show aaa
show ldap

Notes

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. Values: 2 and 3.
Default	3	
Configuration Mode	Config	
History	3.1.0000 3.4.0000	Updated Example
Role	admin	
Example	<pre> switch (config) # ldap version 3 switch (config) # show ldap User base DN : ou=department,dc=example,dc=com User search scope : subtree Login attribute : uid Bind DN : my-dn Bind password : my-password Group base DN : my-group-dn Group attribute : member LDAP version : 3 Referrals : no Server port : 1111 Search Timeout : 5 Bind Timeout : 10 SSL mode : none Server SSL port : 636 (not active) SSL ciphers : TLS1.2 (not active) SSL cert verify : yes SSL ca-list : default-ca-list LDAP servers: 1: 10.10.10.10 2: 10.10.10.12 switch (config) # </pre>	
Related Commands	show aaa show ldap	
Notes		

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
Role	admin
Example	<pre> switch (config) # show ldap User base DN : ou=department,dc=example,dc=com User search scope : subtree Login attribute : uid Bind DN : my-dn Bind password : my-password Group base DN : my-group-dn Group attribute : member LDAP version : 3 Referrals : no Server port : 1111 Search Timeout : 5 Bind Timeout : 10 SSL mode : none Server SSL port : 636 (not active) SSL ciphers : TLS1.2 (not active) SSL cert verify : yes SSL ca-list : default-ca-list LDAP servers: 1: 10.10.10.10 2: 10.10.10.12 switch (config) # </pre>
Related Commands	show aaa show ldap
Notes	

4.13.4.6 System Secure Mode

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
Role	admin
Example	<pre>switch (config) # system secure-mode enable</pre> 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	<pre>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</pre>
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.

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
Role	admin
Example	<pre>switch (config) # show system secure-mode Secure mode configured: yes Secure mode enabled : yes switch (config) #</pre>
Related Commands	system secure-mode enable
Notes	“Secure mode configuration” describes the user configuration “Secure mode enabled” describes the system state

4.14 Cryptographic (X.509, IPSec) and Encryption

This chapter 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.

4.14.1 System File Encryption

This feature encrypts all sensitive data on Mellanox systems including logs certificates, keys, etc.

➤ **To activate encryption on the switch:**

Step 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
- Do not power-off, wait for the system to boot

Type 'YES' to confirm this action: YES
```



*****IMPORTANT NOTE*****

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.

Step 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 MLNX-OS version.



If encryption is enabled, upgrading to a new MLNX-OS® version maintains the encryption configuration.

4.14.1.1 Commands

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	
Role	admin	
Example	<pre>switch (config)# crypto encrypt-data key-location usb key mypassword</pre> Warning! All sensitive files are about to be encrypted - System will perform reset factory, configuration files will be preserved - System will be rebooted - Do not power-off, wait for the system to boot Type 'YES' to confirm this action: YES	
Related Commands		
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.	

crypto ipsec ike

crypto ipsec ike {clear sa [peer {any | <IPv4 or IPv6 address>} local <IPv4 or IPv6 address>] | restart}

Manage 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	
Role	admin	
Example	switch (config)# crypto ipsec ike restart switch (config)#	
Related Commands	N/A	
Notes		

crypto ipsec peer local

```
crypto ipsec peer <IPv4 or IPv6 address> local <IPv4 or IPv6 address> {enable |  
keying {ike [auth {hmac-md5 | hmac-sha1 | hmac-sha256 | null} | dh-group | dis-  
able | encrypt | 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. It has the following optional parameters: • 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 Forward 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

Default	N/A
Configuration Mode	Config
History	3.2.3000
Role	admin
Example	switch (config)# crypto ipsec peer 10.10.10.10 local 10.7.34.139 enable switch (config)#
Related Commands	N/A
Notes	

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	
Role	admin	
Example	switch (config) # crypto certificate default-cert name test	
Related Commands	N/A	
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	

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	
Role	admin	
Example	switch (config) # crypto certificate default-cert name test	
Related Commands	N/A	
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	

crypto certificate generation

crypto certificate generation default {country-code | days-valid | 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 value for days valid.
	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.
Default	N/A	
Configuration Mode	Config	
History	3.2.1000	First version
	3.3.4350	Added “hash-algorithm” parameter
Role	admin	
Example	<pre>switch (config) # crypto certificate generation default hash-algorithm sha256</pre>	

Related Commands	N/A

Notes	The default hashing algorithm used is sha1.

crypto certificate name

```
crypto certificate name {<cert-name> | system-self-signed} {comment <new  
comment> | generate self-signed [comment <cert-comment> | common-name  
<domain> | country-code <code> | days-valid <days> | 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> | 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 - 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
	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	First version
	3.3.4402	Added “hash-algorithm” parameter
Role	admin	
Example	switch (config) # crypto certificate name system-self-signed generate self-signed hash-algorithm sha256	
Related Commands	N/A	
Notes		

crypto certificate system-self-signed

crypto certificate system-self-signed regenerate [days-valid <days>]

Configures default values for certificate generation.

Syntax Description	days-valid	Specifies the number of days the certificate is valid
Default	N/A	
Configuration Mode	Config	
History	3.2.1000	
Role	admin	
Example	switch (config) # crypto certificate system-self-signed regenerate days-valid 3	
Related Commands	N/A	
Notes		

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	
Role	admin	

Example

```
switch (config)# show crypto certificate
Certificate with name 'system-self-signed' (default-cert)
  Comment:                                system-generated self-signed certifi-
icate
  Private Key:                             present
  Serial Number:                           0x546c935511bcafc21ac0e8249fbe0844
  SHA-1 Fingerprint:                       fe6df38dd26801971cb2d44f62d-
be492b6063c5f

  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:

switch (config)#
```

Related Commands

N/A

Notes

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
Role	admin
Example	<pre>switch (config)# show crypto encrypt-data Sensitive files encryption: Status: enabled Key location: usb Cipher: aes256 switch (config)#</pre>
Related Commands	N/A
Notes	

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
Role	admin
Example	<pre> 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. switch (config)# </pre>
Related Commands	N/A
Notes	

4.15 Scheduled Jobs

Use the commands in this section to manage and schedule the execution of jobs

4.15.1 Commands

job

job <job ID>
no job <job ID>

Creates a job.
The no form of the command deletes the job.

Syntax Description	job ID	An integer.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # job 100 switch (config job 100) #	
Related Commands	show jobs	
Notes	Job state is lost on reboot.	

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	
Role	admin	
Example	<pre>switch (config)# job 100 switch (config job 100) # command 10 "show power" switch (config job 100) #</pre>	
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”.	

comment

comment <comment>
no comment

Adds a comment to the job.
The no form of the command deletes the comment.

Syntax Description	comment	The comment to be added (string).
Default	""	
Configuration Mode	Config job	
History	3.1.0000	
Role	admin	
Example	switch (config)# job 100 switch (config job 100) # comment Job_for_example switch (config job 100) #	
Related Commands	show jobs	
Notes		

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
Role	admin
Example	switch (config)# job 100 switch (config job 100) # enable switch (config job 100) #
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.

execute

execute

Forces an immediate execution of the job.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config job
History	3.1.0000
Role	admin
Example	<pre>switch (config)# job 100 switch (config job 100) # execute switch (config job 100) #</pre>
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

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
Role	admin
Example	switch (config)# job 100 switch (config job 100) # fail-continue switch (config job 100) #
Related Commands	show jobs
Notes	

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	“”.	
Configuration Mode	Config job	
History	3.1.0000	
Role	admin	
Example	switch (config)# job 100 switch (config job 100) # name my-job switch (config job 100) #	
Related Commands	show jobs	
Notes		

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	
Role	admin	
Example	<pre>switch (config)# job 100 switch (config job 100) # schedule type once switch (config job 100) #</pre>	
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.	

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	
Role	admin	
Example	<pre>switch (config)# job 100 switch (config job 100) # schedule monthly interval 10 switch (config job 100) #</pre>	
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.	

show jobs

show jobs [<job-id>]

Displays configuration and state (including results of last execution, if any exist) of all jobs, or of one job if a job ID is specified.

Syntax Description	job-id	Job ID.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>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 power Last Output: ===== Module Status ===== PS1 OK PS2 NOT PRESENT switch (config) #</pre>	
Related Commands	show jobs	
Notes		

4.16 Statistics and Alarms

4.16.1 Commands

stats alarm <alarm-id> 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	
Role	admin	
Example	switch (config) # stats alarm cpu_util_indiv clear switch (config) #	
Related Commands	show stats alarm	
Notes		

stats alarm <alarm-id> 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: • <code>cpu_util_indiv</code> - Average CPU utilization too high: percent utilization • <code>disk_io</code> - Operating System Disk I/O per second too high: kilobytes per second • <code>fs_mnt</code> - Free filesystem space too low: percent of disk space free • <code>intf_util</code> - Network utilization too high: bytes per second • <code>memory_pct_used</code> - Too much memory in use: percent of physical memory used • <code>paging</code> - Paging activity too high: page faults • <code>temperature</code> - Temperature is too high: degrees
Default	The default is different per alarm-id	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # stats alarm cpu_util_indiv enable switch (config) #</pre>	
Related Commands	show stats alarm	
Notes		

stats alarm <alarm-id> 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.

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	
Role	monitor/admin	
Example	switch (config) # stats alarm cpu_util_indiv event-repeat single switch (config) #	
Related Commands	show stats alarm	
Notes		

stats alarm <alarm-id> {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	
Role	admin	
Example	switch (config) # stats alarm cpu_util_indiv falling clear-threshold 10 switch (config) #	
Related Commands	show stats alarm	
Notes	Not all alarms support all four thresholds.	

stats alarm <alarm-id> 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	
Role	monitor/admin	
Example	switch (config) # stats alarm paging rate-limit window long 2000 switch (config) #	
Related Commands	show stats alarm	
Notes		

stats chd <chd-id> 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) • eth_day • eth_hour • eth_ip_day • eth_ip_hour • fs_mnt_day - Filesystem system usage average: bytes • fs_mnt_month - Filesystem system usage average: bytes • fs_mnt_week - Filesystem system usage average: bytes • ib_day • ib_hour • 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
Default		N/A	
Configuration Mode		Config	
History		3.1.0000	
Role		admin	

Example

```
switch (config) # stats chd memory_day clear
switch (config) #
```

Related Commands

```
show stats chd
```

Notes

stats chd <chd-id> 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 - eth_day - eth_hour - fs_mnt_day - Filesystem system usage average: bytes - fs_mnt_month - Filesystem system usage average: bytes - fs_mnt_week - Filesystem system usage average: bytes - ib_day - ib_hour - 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
Default	Enabled	
Configuration Mode	Config	
History	3.1.0000	
Role	monitor/admin	

Example	switch (config) # stats chd memory_day enable switch (config) #

Related Commands	show stats chd

Notes	

stats chd <chd-id> 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 • eth_day • eth_hour • fs_mnt_day - Filesystem system usage average: bytes • fs_mnt_month - Filesystem system usage average: bytes • fs_mnt_week - Filesystem system usage average: bytes • ib_day • ib_hour • 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
	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
Role	monitor/admin
Example	<pre>switch (config) # stats chd memory_day compute time interval 120 switch (config) # show stats chd memory_day CHD "memory_day" (Average physical memory usage: bytes): Source dataset: sample "memory" Computation basis: time Interval: 120 second(s) Range: 1800 second(s) switch (config) #</pre>
Related Commands	show stats chd
Notes	

stats sample <sample-id> 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 • eth • eth-abs • eth_ip • fan - Fan speed • fs_mnt_bytes - Filesystem usage: bytes • fs_mnt_inodes - Filesystem usage: inodes • ib • 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
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # stats sample temperature clear switch (config) #</pre>	
Related Commands	show stats sample	
Notes		

stats sample <sample-id> enable

stats sample <sample-id> enable
no stats 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 • eth • fan - Fan speed • fs_mnt_bytes - Filesystem usage: bytes • fs_mnt_inodes - Filesystem usage: inodes • ib • 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
Default	Enabled	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # stats sample temperature enable switch (config) #	
Related Commands	show stats sample	
Notes		

stats sample <sample-id> interval

stats sample <sample ID> interval <number of seconds>

Sets the amount of time between samples for the specified group of sample data.

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 • eth • fan - Fan speed • fs_mnt_bytes - Filesystem usage: bytes • fs_mnt_inodes - Filesystem usage: inodes • ib • 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
	number of seconds	Interval in seconds.
Default	Different per sample	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # stats sample temperature interval 1 switch (config) # show stats sample temperature Sample "temperature" (Modules temperature): Enabled: yes Sampling interval: 1 second switch (config) #</pre>	
Related Commands	show stats sample	
Notes		

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
Role	admin
Example	switch (config) # stats clear-all switch (config) #
Related Commands	N/A
Notes	

stats export

stats export <format> <report name> [{after | before} <yyyy/mm/dd>
<hh:mm:ss>] [filename <filename>]

Exports statistics to a file.

Syntax Description	format	Currently the only supported value for <format> is “csv” (comma-separated value).
	report name	Determines dataset to be exported. Possible report names are: - memory - Memory utilization - paging - Paging I/O - cpu_util - CPU utilization
	after before	Only includes stats collected after or before a specific time.
	yyyy/mm/dd	Date: It must be between 1970/01/01 and 2038/01/19.
	hh:mm:ss	Time: It must be between 00:00:00 and 03:14:07 UTC and is treated as local time.
	filename	Specifies filename to give new report. If a filename is specified, the stats will be exported to a file of that name; otherwise a name will be chosen automatically and will contain the name of the report and the time and date of the export. Any automatically-chosen name will be given a .csv extension.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # stats export csv memory filename mellanoxexample before 2000/08/14 15:59:50 after 2000/08/14 15:01:50 Generated report file: mellanoxexample.csv switch (config) # show files stats mellanoxexample.csv switch (config) #</pre>	

Related Commands show files stats

Notes

show stats alarm

show stats alarm [<Alarm ID> [rate-limit]]

Displays status of all alarms or the specified alarm.

Syntax Description	Alarm ID	May be: • <code>cpu_util_indiv</code> - Average CPU utilization too high: percent utilization • <code>disk_io</code> - Operating System Disk I/O per second too high: kilobytes per second • <code>fs_mnt</code> - Free filesystem space too low: percent of disk space free • <code>intf_util</code> - Network utilization too high: bytes per second • <code>memory_pct_used</code> - Too much memory in use: percent of physical memory used • <code>paging</code> - Paging activity too high: page faults • <code>temperature</code> - Temperature is too high: degrees
	rate-limit	Displays rate limit parameters.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>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): (dis- abled) 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 switch (config) #</pre>	
Related Commands	stats alarm	
Notes		

show stats chd

show stats chd [<CHD ID>]

Displays configuration of all statistics CHDs.

Syntax Description	CHD ID	May be: • <code>cpu_util_indiv</code> - Average CPU utilization too high: percent utilization • <code>disk_io</code> - Operating System Disk I/O per second too high: kilobytes per second • <code>fs_mnt</code> - Free filesystem space too low: percent of disk space free • <code>intf_util</code> - Network utilization too high: bytes per second • <code>memory_pct_used</code> - Too much memory in use: percent of physical memory used • <code>paging</code> - Paging activity too high: page faults • <code>temperature</code> - Temperature is too high: degrees
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>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) switch (config) #</pre>	
Related Commands	stats chd	
Notes		

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	Config
History	3.1.0000
Role	admin
Example	<pre>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% switch (config) #</pre>
Related Commands	N/A
Notes	

show stats sample

show stats sample [<sample ID>]

Displays sampling interval for all samples, or the specified one.

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 • eth • fan - Fan speed • fs_mnt_bytes - Filesystem usage: bytes • fs_mnt_inodes - Filesystem usage: inodes • ib • 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
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # show stats sample fan Sample "fan" (Fan speed): Enabled: yes Sampling interval: 1 minute 11 seconds switch (config) #</pre>	
Related Commands	N/A	
Notes		

4.17 Chassis Management

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

Table 30 - Chassis Manager 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

4.17.1 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 (leaf, spine, 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

4.17.1.1 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.

4.17.1.2 System Health Monitor Alerts Scenarios

- System Health Monitor sends notification alerts in the following cases:

Table 31 - System Health Monitor Alerts Scenarios (Sheet 1 of 7)

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 <fan_number> speed in spine number <spine_number> is below minimal range	A spine fan speed is below minimal threshold: 30% of maximum speed	Email, fan LED and system status LED set red, log alert, SNMP	Check the fan and replace it if required	"Fan speed <fan_number> in spine number <spine_number> has been restored to its normal state"
<fan_name> is unresponsive	A chassis fan is not responsive on MLNX-OS systems	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 <fan_number> in spine number <spine_number> is unresponsive	A spine fan is not responsive on MLNX-OS systems	Email, fan LED and system status LED set red, log alert, SNMP	Check fan connectivity and replace it if required	"Fan <fan_number> in spine number <spine_number> 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"
Fan <fan_number> in spine number <spine_number> is not present.	A spine fan is missing	Email, fan LED and system status LED set red, log alert, SNMP	Insert a fan unit	"Fan <fan_number> in spine number <spine_number> 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"

Table 31 - System Health Monitor Alerts Scenarios (Sheet 2 of 7)

Alert Message	Scenario	Notification Indicator	Recovery Action	Recovery Message
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 MLNX-OS systems	Email, power supply LED and system status LED set red, log alert, SNMP	Check chassis fans connections. On MLNX-OS systems, check system fan connections.	“Power supply <ps_number> temperature is back to normal”
Power Supply <number> is unresponsive	A power supply is malfunctioning or disconnected	Email, system status 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”
Unit/leaf/spine <leaf/spine number> is unresponsive	A leaf/spine is not responsive	Email, system status LED set red, log alert, SNMP	Check leaf/spine connectivity and replace it if required	“Leaf/spine number <leaf/spine number> has been restored to its normal state”
Unit/leaf/spine voltage is out of range	One of the voltages in a MLNX-OS unit is below minimal threshold or higher than the maximum threshold - both thresholds are 15% of the expected voltage	Email, system status LED set red, log alert, SNMP	Check leaf connectivity	“Unit voltage is in range”
ASIC temperature is too hot	A SwitchX unit temperature is higher than the maximum threshold of 105 Celsius on MLNX-OS systems	Email, system status LED set red, log alert, SNMP	Check the fans system	“SwitchX temperature is back to normal”

Table 31 - System Health Monitor Alerts Scenarios (Sheet 3 of 7)

Alert Message	Scenario	Notification Indicator	Recovery Action	Recovery Message
BBU Health Monitoring				
“BBU<num> active alarms: Under-tem- perature during discharge (UTD)”	Under-temperature during discharge	Email, system sta- tus LED set red, log alert, SNMP	Check ambi- ent tempera- ture. Replace BBU if the problem per- sist.	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Under-tem- perature during charge (UTC)”	Under-temperature during charge	Email, system sta- tus LED set red, log alert, SNMP	Check ambi- ent tempera- ture. Replace BBU if the problem per- sist.	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Over pre- charge current (PCHGC)”	Over pre-charge current	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU if the prob- lem persists.	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Overcharging voltage (CHGV)”	Overcharging voltage	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU if the prob- lem persists.	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Overcharging current (CHGC)”	Overcharging current	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU if the prob- lem persists.	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Overcharge (OC)”	Overcharged BBU	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU if the prob- lem persists.	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Charge time- out suspend (CTOS)”	Charge timeout suspend	Email, system sta- tus LED set red, log alert, SNMP	N/A	“Module BBU<num> has been restored to its normal state”

Table 31 - System Health Monitor Alerts Scenarios (Sheet 4 of 7)

Alert Message	Scenario	Notification Indicator	Recovery Action	Recovery Message
“BBU<num> active alarms: Charge time- out (CTO)”	Charge timeout	Email, system sta- tus LED set red, log alert, SNMP	N/A	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Pre-charge tim- eout suspend (PTOS)”	Pre-charge timeout sus- pend	Email, system sta- tus LED set red, log alert, SNMP	N/A	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Pre-charge tim- eout (PTO)”	Pre-charge timeout	Email, system sta- tus LED set red, log alert, SNMP	N/A	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Over-tempera- ture FET (OTF)”	Over-temperature FET	Email, system sta- tus LED set red, log alert, SNMP	N/A	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Cell under- voltage com- pensated (CUVC)”	Cell under-voltage com- pensated	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU if the prob- lem persists.	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Over-tempera- ture during dis- charge (OTD)”	Over-temperature during discharge	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU if the prob- lem persists.	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Over-tempera- ture during charge (OTC)”	Over-temperature during charge	Email, system sta- tus LED set red, log alert, SNMP	Check ambi- ent tempera- ture. Replace BBU if the problem per- sists.	“Module BBU<num> has been restored to its normal state”

Table 31 - System Health Monitor Alerts Scenarios (Sheet 5 of 7)

Alert Message	Scenario	Notification Indicator	Recovery Action	Recovery Message
“BBU<num> active alarms: Short-circuit during dis- charge latch (ASCDL)”	Short-circuit during dis- charge latch	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU	N/A
“BBU<num> active alarms: Short-circuit during dis- charge (ASCL)”	Short-circuit during dis- charge	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU if the prob- lem persists.	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Short-circuit during charge latch (ASCCL)”	Short-circuit during charge latch	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU	N/A
“BBU<num> active alarms: Short-circuit during charge (ASCC)”	Short-circuit during charge	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU if the prob- lem persists.	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Overload during dis- charge latch (AOLDL)”	Overload during dis- charge latch	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU	N/A
“BBU<num> active alarms: Overload during dis- charge (AOLD)”	Overload during dis- charge	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU if the prob- lem persists.	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Over-current during dis- charge 1 (OCD1)”	Over-current during dis- charge	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU if the prob- lem persists.	“Module BBU<num> has been restored to its normal state”

Table 31 - System Health Monitor Alerts Scenarios (Sheet 6 of 7)

Alert Message	Scenario	Notification Indicator	Recovery Action	Recovery Message
“BBU<num> active alarms: Over-current during dis- charge 2 (OCD2)”	Over-current during dis- charge	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU if the prob- lem persists.	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Over-current during charge 1 (OCC1)”	Over-current during charge	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU if the prob- lem persists.	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Over-current during charge 2 (OCC2)”	Over-current during charge	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU if the prob- lem persists.	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Cell over-volt- age (COV)”	Cell over-voltage	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU if the prob- lem persists (may take up to 48 hours if BBU was in storage).	“Module BBU<num> has been restored to its normal state”
“BBU<num> active alarms: Cell under- voltage (CUV)”	Cell under-voltage	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU if the prob- lem persists.	“Module BBU<num> has been restored to its normal state”
“Module BBU<num> voltage is out of range”	Cell over-voltage	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU if the prob- lem persists (may take up to 48 hours if BBU was in storage).	“Module BBU<num> voltage is back in range”
“Module BBU<num> current is too high”	Over-current during charge	Email, system sta- tus LED set red, log alert, SNMP	Replace BBU if the prob- lem persists.	“BBU<num> has been restored to its normal state”

Table 31 - System Health Monitor Alerts Scenarios (Sheet 7 of 7)

Alert Message	Scenario	Notification Indicator	Recovery Action	Recovery Message
“Module BBU<num> current is too high”	Over-current during discharge	Email, system status LED set red, log alert, SNMP	Replace BBU if the problem persists.	“BBU<num> has been restored to its normal state”
“Module BBU<num> temperature is too hot”	Over-temperature during charge	Email, system status LED set red, log alert, SNMP	Check ambient temperature. Replace BBU if the problem persists.	“Module BBU<num> temperature is back to normal”
“Module BBU<num> temperature is too hot”	Over-temperature during discharge	Email, system status LED set red, log alert, SNMP	Check ambient temperature. Replace BBU if the problem persists.	“Module BBU<num> temperature is back to normal”

4.17.2 Power Management

4.17.2.1 Power Supply Options

MLNX-OS offers power redundancy configurations and monitoring for director switch systems. Director switch systems have the following redundancy configuration modes:

- “combined” – no power supply is reserved, the redundancy is not enabled.
- “ps-redundant” – one power supply unit is redundant to the rest. The system can work with one less power supply unit.
- “grid-redundant” – the power supplies are split into two logical power supply grids, first half of the PSUs belongs to grid A and the second half to grid B. The systems can work with only one grid. When using grid-redundancy mode the power budget is calculated according to the minimum power budget between the grids. This mode is available only in SX65xx-NR chassis systems. During switch initialization, or hot-plugging of switch components, MLNX-OS enables and/or disables switch components according to the available power budget.

MLNX-OS may send power alarms (via SNMP or email) as follow:

- If the available budget is insufficient for all the system components an `insufficientPower` event is generated. In this mode several switch components may be disabled.
- If the total power of the system is insufficient for redundancy, a `lowPower` event is generated.

- If a connected power supply provides below 1.6K Watts or grid-redundancy mode is configured and a power supply is connected to a 110V grid, then a powerRedundancy-Mismatch event is generated, where grid redundancy can not be achieved in such configuration.

In case of an insufficient-power mode, the order in which the FRUs are turned ON is first spines (1,2,3...max) and then the leafs (1,2,3...max), while the order of the FRUs in case of turning them OFF is first the spines (max...3) and then the leafs (max...1). The management modules are not affected.

For the trap OID, please refer to the Mellanox-MIB file.



Power cycle is needed after changing power redundancy mode on a director switch system.

4.17.2.2 Width Reduction Power Saving

Link width reduction (LWR) is a Mellanox 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 Mellanox 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 and InfiniBand FDR 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.



See command “power-management width” on page 517.

Table 32 - LWR Configuration Behavior

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

Table 32 - LWR Configuration Behavior

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

4.17.2.3 Managing Chassis Power

It is possible to shut down or power up modules in a chassis by using the `power enable` and `no power enable` commands.

Step 1. Change to Config mod. Run:

```
switch [standalone: master] > enable
switch [standalone: master] # configure terminal
switch [standalone: master] (config) #
```

Step 2. Run the command `show power` to get a list of modules that are available to power up or down.

Step 3. To power down a module run the command `no power enable` followed by a module.

```
switch [standalone: master] (config) # no power enable ps1
```

Step 4. To power up a module run the command `power enable` followed by a module.

```
switch [standalone: master] (config) # power enable ps1
```

Using the `show power` command it is possible to see the power consumption of the system and also the power consumption by power supply unit.

4.17.3 Monitoring Environmental Conditions

Step 1. Display module's temperature. Run:

```
switch (config) # show temperature
=====
Module  Sensor                CurTemp  Status
                (Celsius)
=====
MGMT    CPU_BOARD_MONITOR      40.00    OK
L01     BOARD_MONITOR          27.00    OK
L01     QSFP_TEMP1             24.00    OK
L01     QSFP_TEMP2             22.00    OK
L01     QSFP_TEMP3             21.00    OK
L01     SX                     38.00    OK
L02     BOARD_MONITOR          27.00    OK
L02     QSFP_TEMP1             24.50    OK
L02     QSFP_TEMP2             22.50    OK
L02     QSFP_TEMP3             21.50    OK
L02     SX                     32.00    OK
PS2     PS_MONITOR             24.66    OK
PS3     PS_MONITOR             31.04    OK
PS4     PS_MONITOR             28.06    OK
S01     BOARD_MONITOR          23.00    OK
S01     SX                     34.00    OK
S01     SX_AMBIENT_TEMP        22.50    OK
S02     BOARD_MONITOR          24.00    OK
S02     SX                     49.00    OK
S02     SX_AMBIENT_TEMP        24.00    OK
switch (config) #
```


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

```
switch (config) # show voltage
=====
Module  Power Meter      Reg Expected Actual  Status
        Voltage      Voltage
=====
PS2     PS_MONITOR        V1  48.00   46.88   OK
PS3     PS_MONITOR        V1  48.00   48.29   OK
PS4     PS_MONITOR        V1  48.00   48.29   OK
MGMT    CPU_BOARD_MONITOR V1  12.00   11.92   OK
MGMT    CPU_BOARD_MONITOR V2   2.50    2.48   OK
MGMT    CPU_BOARD_MONITOR V3   3.30    3.31   OK
MGMT    CPU_BOARD_MONITOR V4   3.30    3.30   OK
MGMT    CPU_BOARD_MONITOR V5   1.80    1.81   OK
MGMT    CPU_BOARD_MONITOR V6   1.20    1.26   OK
S01     BOARD_MONITOR     V1   3.30    3.33   OK
S01     BOARD_MONITOR     V2   2.27    2.15   OK
S01     BOARD_MONITOR     V3   1.80    1.76   OK
S01     BOARD_MONITOR     V4   3.30    3.30   OK
S01     BOARD_MONITOR     V5   0.90    0.93   OK
S01     BOARD_MONITOR     V6   1.20    1.19   OK
S02     BOARD_MONITOR     V1   3.30    3.26   OK
S02     BOARD_MONITOR     V2   2.27    2.16   OK
S02     BOARD_MONITOR     V3   1.80    1.79   OK
S02     BOARD_MONITOR     V4   3.30    3.31   OK
S02     BOARD_MONITOR     V5   0.90    0.95   OK
S02     BOARD_MONITOR     V6   1.20    1.20   OK
L01     BOARD_MONITOR     V1   3.30    3.33   OK
L01     BOARD_MONITOR     V2   2.27    2.16   OK
L01     BOARD_MONITOR     V3   1.80    1.76   OK
L01     BOARD_MONITOR     V4   3.30    3.30   OK
L01     BOARD_MONITOR     V5   0.90    0.93   OK
L01     BOARD_MONITOR     V6   1.20    1.19   OK
L02     BOARD_MONITOR     V1   3.30    3.26   OK
L02     BOARD_MONITOR     V2   2.27    2.17   OK
L02     BOARD_MONITOR     V3   1.80    1.79   OK
L02     BOARD_MONITOR     V4   3.30    3.30   OK
L02     BOARD_MONITOR     V5   0.90    0.89   OK
L02     BOARD_MONITOR     V6   1.20    1.19   OK
switch (config) #
```

Step 3. Display the fan speed and status. Run:

```
switch (config) # show fan
=====
Module           Device      Fan  Speed      Status
              (RPM)
=====
FAN1             FAN        F1   6994.00    OK
FAN2             FAN        F1   6792.00    OK
FAN3             FAN        F1   6870.00    OK
FAN4             FAN        F1   6818.00    OK
S01              FAN        F1   7800.00    OK
S01              FAN        F2   8130.00    OK
S02              FAN        F1   8130.00    OK
S02              FAN        F2   8490.00    OK
S03              FAN        -    -          NOT PRESENT
S04              FAN        -    -          NOT PRESENT
S05              FAN        -    -          NOT PRESENT
S06              FAN        -    -          NOT PRESENT
switch (config) #
```

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

```
switch (config) # show power consumers
=====
Module           Power    Voltage    Current    Status
              (Watts)
=====
FAN1             15.55    48.00     0.32      OK
FAN2             16.26    48.00     0.34      OK
FAN3             15.30    48.00     0.32      OK
FAN4             14.98    48.00     0.31      OK
L01              32.45    48.00     0.68      OK
L02              28.75    48.00     0.60      OK
MGMT             16.08    48.00     0.34      OK
S01              37.34    48.00     0.78      OK
S02              35.09    48.00     0.73      OK

Total power used : 211.79 W
Max power : 686.00 W
switch (config) #
```

4.17.4 USB Access

MLNX-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://admin:admin@127.0.0.1/var/mnt/usb1/image.img
```

- **To save log file 'my-logfile' to a USB device under the name test_logfile using the logging files command, run (in Enable or Config mode):**

```
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 (in Enable or Config mode).**

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

4.17.5 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    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
```

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

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

4.17.6 High Availability (HA)

Mellanox high end management director switch systems support redundant management modules. Chassis HA reduces downtime as it assures continuity of the work even when a management module dies. Chassis HA management allows the systems administrator to associate a single IP address with the appliance. Connecting to that IP address allows the user to change and review the system's chassis parameters regardless of the active management module.

Figure 16: SX65xx with Dual Management Modules



4.17.6.1 Chassis High Availability Nodes Roles

Every node in the Chassis HA has one of the following roles/modes:

- Master – the node that manages chassis configurations and services the chassis IP addresses
- Slave – the node that replaces the Master node and takes over its responsibilities once the Master node is down.



The master node is the only node that has access to chassis components such as temperature, inventory and firmware.

The CPU role of the current management node can be recognized by the following methods:

- Run the `show chassis ha` command.

```
switch (config) # show chassis ha
2-node HA state:
  Box management IP: 172.30.1.200/16
    interface: mgmt0

  local role: master
  local slot: 1
  other state: ready
  reset count: 0

switch (config) #
```

- Check the LEDs in the management modules as displayed in the figure below.

Figure 17: SX60xx's LEDs



- Go to the WebUI => System => Modules page and see the information on the LEDs.

4.17.6.2 Malfunctioned CPU Behavior

When a CPU is not responding to an internal communication with the other CPU, the non responding CPU will be reset by the other CPU. Each time a CPU resets, a counter is incremented. After 5 resets a CPU is considered malfunctioned and will be shut down.

To verify how many times a CPU is reset, run the following command:

```
switch-11a14e [default: master] (config) # show chassis ha
2-node HA state:
  Box management IP: 172.30.1.200/16
    interface: mgmt0

    local role: master
    local slot: 1
    other state: ready
    reset count: 1

switch-11a14e [default: master] (config) #
```

To verify if a CPU has been shut down, either run the following command:

```
switch-11a14e [default: master] (config) # show chassis ha
2-node HA state:
  Box management IP: 172.30.1.200/16
    interface: mgmt0

    local role: master
    local slot: 1
    other state: powered-off
    reset count: 5

switch-11a14e [default: master] (config) #
```

Or check the system page in the WebUI, the management figure will be grayed out.

To enable the malfunctioned CPU, first replace it and run `chassis ha reset other`.

4.17.6.3 Box IP Centralized Location

Box IP (BIP) centralized management infrastructure enables you to configure and monitor the system. The BIP continues to function even if one of the management blades dies. Box IP is defined by running the `chassis ha bip <board IP address>` command. The created BIP is used as the master IP's alias.

Example:

```
SX648 [standalone: master] (config) # chassis ha bip 192.168.10.100 255.255.255.0
SX648 [standalone: master] (config) #
```

4.17.6.4 System Configuration

System configuration changes should be performed by the master using the BIP otherwise they are overridden by the master configuration.

Chassis HA is based on database replication enabling the entire master configuration to be replicated to the slave. Data such as chassis configuration is replicated. However, run time information such as time, logs, active user lists, is not copied. Additionally, node specific configuration information such as host name and IP address is not copied..



Chassis HA requires connectivity of both management modules (mgmt0, mgmt1) in the same broadcast domain.

4.17.6.5 Takeover Functionally

Management CPU functional takeover takes up to 20-30 seconds. However, when plugging in a module, you need to wait for approximately 3 minutes before making any other hardware change. During the takeover process, the Master LED status is differentiated by a color scheme. To verify the system's status, run the "show chassis ha" command on both managements.

In case of CPU malfunction the system tries to reset it 5 times to solve the issue. If the CPU is not activated after resetting, the system powers it off as well as its attached spine. Once the CPU is powered off, the user should replace the malfunctioned CPU module. To power on the CPU and the attached spine, plug the module in, log into the Master CPU and run the "chassis ha power enable other" command.



Although the LEDs are functional during the takeover, wait for approximately 3 minutes before making any other hardware change.

Master example

```
switch [default: master] (config) # show chassis ha
2-node HA state:
  Box management IP: 172.30.1.200/16
    interface: mgmt0

    local role: master
    local slot: 1
    other state: ready
    reset count: 1

switch [default: master] (config) #
```

Slave example

```
switch [default: master] (config) # show chassis ha
2-node HA state:
  Box management IP: 172.30.1.200/16
  interface: mgmt0

  local role: slave
  local slot: 2
  other state: ready
  reset count: 0

switch [default: master] (config) #
```



Not following these instructions may result in some errors in the log. These errors may be safely ignored.

4.17.7 System Reboot

4.17.7.1 Rebooting 1U Switches

➤ *To reboot a 1U switch system:*

Step 1. Enter Enable or Config mode. Run:

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

Step 2. Reboot the system. Run:

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

4.17.7.2 Rebooting Director Switches

Mellanox high end management director switch systems support redundant management modules. Chassis HA reduces downtime as it assures continuity of the work even when a management module dies. Chassis HA management allows the systems administrator to associate a single IP address with the appliance. Connecting to that IP address allows the user to change and review the system's chassis parameters regardless of the active management module.

➤ *To reboot director switches:*

Step 1. Connect to BIP. Please refer to [Section 4.17.6.3, "Box IP Centralized Location," on page 509](#) for more information.

Step 2. Enter Config mode. Run:

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

Step 3. Reboot the slave management. Run:

```
switch [default: master] (config) # chassis ha reset other  
switch [default: master] (config) #
```

Step 4. Reboot the master management. Run:

```
switch [default: master] (config) # reload
```


4.17.8 Commands

4.17.8.1 Chassis Management

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 Interface Port Channel	
History	3.2.3000	
Role	admin	
Example	switch (config) # clear counters	
Related Commands		
Notes		

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 seconds.
	report-clear	Clears the health report.
Default	max-report-len: 50 re-notif-cntr:	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # health re-notif-cntr 125 switch (config) #	
Related Commands	show health-report	
Notes		

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	Added director switch support
Role	admin	
Example	switch (config) # led MGMT uid on switch (config) #	
Related Commands	N/A	
Notes	- On 1U switch systems, the module parameter can only be MGMT - On director switch systems, the module parameter may be MGMT#, L#, S# (e.g. MGMT1, L01, S01)	

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	
Role	admin	
Example	switch (config) # power enable L01 switch (config) #	
Related Commands	show power show power consumers	
Notes	• This command is not applicable on 1U systems • It is recommended to run this command prior to extracting a module from the switch system, else errors are printed in the log	

power-management width

power-management width {auto | force}
no power-management width

Sets the width of the interface to be automatically adjusted.
The no form of the command disables power-saving.

Syntax Description	auto	Allows the system to automatically decide whether to work in power-saving mode or not.
	force	Forces power-saving mode on the port.
Default	Disabled	
Configuration Mode	Config Interface IB Config Interface Ethernet	
History	3.3.4000	
Role	admin	
Example	switch (config interface ib 1/1) # power-management width auto switch (config) #	
Related Commands	show interface	
Notes		

power redundancy-mode

power redundancy-mode [combined | grid-redundant | ps-redundant]

Controls the power supply redundancy mode.

Syntax Description	combined	No redundancy - no alarm threshold.
	grid-redundant	N+N – the alarm threshold will be set to a level, indicating when the power availability falls below power that can support N+N scheme
	ps-redundant	N+1 – the alarm threshold will be set to a level, indicating when the power availability falls below power that can support N+1 scheme
Default	N/A	
Configuration Mode	Config	
History	3.1.0000, 3.2.0000	
Role	admin	
Example	switch (config) # power redundancy-mode combined switch (config) #	
Related Commands	N/A	
Notes	- The difference between the modes sets the threshold for power supply redundancy failure. It does not change any power supply configuration. - This command is not applicable for 1U or blade systems.	

usb eject

usb eject

Gracefully turns off the USB interface.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # usb eject switch (config) #
Related Commands	N/A
Notes	Applicable only for systems with USB interface.

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 3.4.3050 Updated Example 3.6.1002 Updated Example
Role	admin
Example	<pre>switch (config) # show asic-version ===== Module Device Version ===== MGMT SX 9.2.9160 On a Switch-IB device: switch (config) # show asic-version ===== Module Device Version ===== L05 SIB2-1 15.0200.0092 L05 SIB2-2 15.0200.0092 L06 SIB-1 15.0200.0092 L06 SIB-2 15.0200.0092 switch (config) #</pre>
Related Commands	N/A
Notes	

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
Role	admin
Example	<pre>switch (config) # show bios BIOS version : 4.6.5 BIOS subversion : Official AMI Release BIOS release date : 07/02/2013 switch (config) #</pre>
Related Commands	
Notes	The command is available only on x86 systems

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
Role	admin
Example	<pre>switch (config) # show cpld ===== Name Type Version ===== Cpld1 CPLD_TOR 4 Cpld2 CPLD_PORT1 2 Cpld3 CPLD_PORT2 2 Cpld4 CPLD_MEZZ 3 switch (config) #</pre>
Related Commands	N/A
Notes	

show fan

show fan

Displays fans status.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.1.0000
Role	admin
Example	<pre>switch (config) # show fan switch (config) # show fan ===== Module Device Fan Speed Status (RPM) ===== FAN FAN F1 5340.00 OK FAN FAN F2 5340.00 OK FAN FAN F3 5640.00 OK FAN FAN F4 5640.00 OK PS1 FAN F1 5730.00 OK PS2 FAN - - NOT PRESENT switch (config) #</pre>
Related Commands	N/A
Notes	

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	First version
	3.3.0000	Output update
Role	admin	
Example	<pre>switch (config) # show health-report ===== ALERTS CONFIGURATION ===== Re-notification counter (sec):[3600] Report max counter: [50] ===== HEALTH REPORT ===== No Health issues file switch (config) #</pre>	
Related Commands	N/A	
Notes	Problems with the power supply cannot be monitored on SX1016 switch systems.	

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 output		
Role	admin				
Example	switch (config) # show inventory				

	Module	Part number	Serial Number	Asic Rev.	HW Rev.

	CHASSIS	MSX1036B-1SFR	MT1205X01549	N/A	A1
	MGMT	MSX1036B-1SFR	MT1205X01549	0	A1
	FAN	MSX60-FF	MT1206X07209	N/A	A3
	PS1	MSX60-PF	MT1206X06697	N/A	A2
switch (config) #					
Related Commands	N/A				
Notes					

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	Config																																																																																		
History	3.6.1002																																																																																		
	3.6.2002	Updated output																																																																																	
Role	admin																																																																																		
Example	<pre>switch (config) # show leds Module LED Status -----</pre> <table> <tbody> <tr><td>MGMT1</td><td>STATUS</td><td>Green</td></tr> <tr><td>MGMT1</td><td>REAR_FAN</td><td>Green</td></tr> <tr><td>MGMT1</td><td>PS</td><td>Green</td></tr> <tr><td>MGMT1</td><td>FRONT_FAN</td><td>Green</td></tr> <tr><td>MGMT1</td><td>MASTER/SLAVE</td><td>Green</td></tr> <tr><td>L01</td><td>STATUS</td><td>Green</td></tr> <tr><td>L01</td><td>UID</td><td>Blue</td></tr> <tr><td>L02</td><td>STATUS</td><td>Green</td></tr> <tr><td>L02</td><td>UID</td><td>Blue</td></tr> <tr><td>L03</td><td>STATUS</td><td>Green</td></tr> <tr><td>L03</td><td>UID</td><td>Off</td></tr> <tr><td>L04</td><td>STATUS</td><td>Green</td></tr> <tr><td>L04</td><td>UID</td><td>Off</td></tr> <tr><td>L05</td><td>STATUS</td><td>Green</td></tr> <tr><td>L05</td><td>UID</td><td>Off</td></tr> <tr><td>L06</td><td>STATUS</td><td>Green</td></tr> <tr><td>L06</td><td>UID</td><td>Off</td></tr> <tr><td>S01</td><td>STATUS</td><td>Green</td></tr> <tr><td>S01</td><td>FAN</td><td>Green</td></tr> <tr><td>S02</td><td>STATUS</td><td>Green</td></tr> <tr><td>S02</td><td>FAN</td><td>Green</td></tr> <tr><td>S03</td><td>STATUS</td><td>Green</td></tr> <tr><td>S03</td><td>FAN</td><td>Green</td></tr> <tr><td>FAN1</td><td>STATUS</td><td>Green</td></tr> <tr><td>FAN2</td><td>STATUS</td><td>Green</td></tr> <tr><td>FAN3</td><td>STATUS</td><td>Green</td></tr> <tr><td>FAN4</td><td>STATUS</td><td>Green</td></tr> </tbody> </table>		MGMT1	STATUS	Green	MGMT1	REAR_FAN	Green	MGMT1	PS	Green	MGMT1	FRONT_FAN	Green	MGMT1	MASTER/SLAVE	Green	L01	STATUS	Green	L01	UID	Blue	L02	STATUS	Green	L02	UID	Blue	L03	STATUS	Green	L03	UID	Off	L04	STATUS	Green	L04	UID	Off	L05	STATUS	Green	L05	UID	Off	L06	STATUS	Green	L06	UID	Off	S01	STATUS	Green	S01	FAN	Green	S02	STATUS	Green	S02	FAN	Green	S03	STATUS	Green	S03	FAN	Green	FAN1	STATUS	Green	FAN2	STATUS	Green	FAN3	STATUS	Green	FAN4	STATUS	Green
MGMT1	STATUS	Green																																																																																	
MGMT1	REAR_FAN	Green																																																																																	
MGMT1	PS	Green																																																																																	
MGMT1	FRONT_FAN	Green																																																																																	
MGMT1	MASTER/SLAVE	Green																																																																																	
L01	STATUS	Green																																																																																	
L01	UID	Blue																																																																																	
L02	STATUS	Green																																																																																	
L02	UID	Blue																																																																																	
L03	STATUS	Green																																																																																	
L03	UID	Off																																																																																	
L04	STATUS	Green																																																																																	
L04	UID	Off																																																																																	
L05	STATUS	Green																																																																																	
L05	UID	Off																																																																																	
L06	STATUS	Green																																																																																	
L06	UID	Off																																																																																	
S01	STATUS	Green																																																																																	
S01	FAN	Green																																																																																	
S02	STATUS	Green																																																																																	
S02	FAN	Green																																																																																	
S03	STATUS	Green																																																																																	
S03	FAN	Green																																																																																	
FAN1	STATUS	Green																																																																																	
FAN2	STATUS	Green																																																																																	
FAN3	STATUS	Green																																																																																	
FAN4	STATUS	Green																																																																																	
Related Commands	N/A																																																																																		
Notes																																																																																			

show memory

show memory

Displays memory status.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.1.0000
Role	admin
Example	<pre>switch (config) # show memory Total Used Free Used+B/C Free-B/C Physical 2027 MB 761 MB 1266 MB 1214 MB 813 MB Swap 0 MB 0 MB 0 MB Physical Memory Borrowed for System Buffers and Cache: Buffers: 0 MB Cache: 452 MB Total Buffers/Cache: 452 MB switch (config) #</pre>
Related Commands	N/A
Notes	

show module

show module

Displays modules status.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.1.0000 First version 3.3.0000 Added “Is Fatal” column 3.4.2008 Updated command output 3.4.3000 Updated command output and added note
Role	admin
Example	<pre>switch (config) # show module ===== Module Status ===== MGMT ready FAN1 ready FAN2 ready PS1 ready PS2 not-present switch (config) #</pre>
Related Commands	N/A
Notes	The Status column may have one of the following values: error, fatal, not-present, powered-off, powered-on, ready.

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								
Role	admin								
Example	switch (config) # show power								

	Module	Device	Sensor	Power [Watts]	Voltage [Volts]	Current [Amp]	Capacity [Watts]	Feed	Status

	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	
	switch (config) #								
Related Commands	N/A								
Notes									

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
Role	admin
Example	<pre>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 Total power used : 52.96 Watts switch (config) #</pre>
Related Commands	N/A
Notes	

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
Role	admin	
Example	<pre>switch (config) # show protocols Ethernet spanning-tree enabled lacp rst lldp disabled igmp-snooping disabled ets enabled priority-flow-control disabled sflow disabled openflow disabled mlag disabled dot1x disabled isolation-group disabled IP routing bgp disabled pim disabled vrrp disabled ospf disabled magg disabled dhcp-relay disabled Infiniband sm enabled enabled</pre>	
Related Commands	N/A	
Notes		

show resources

show resources

Displays system resources.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.1.0000
Role	admin
Example	<pre>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% switch (config) #</pre>
Related Commands	N/A
Notes	

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	First version
	3.3.0000	Added gateway support
	3.6.1002	Updated output
Role	admin	
Example	switch (config) # show system capabilities IB: Supported, L2, Adaptive Routing Max SM nodes: 648 IB Max licensed speed: EDR switch (config) # show system capabilities Ethernet: Supported, L2, L3 Ethernet Max licensed speed: 56Gb	
Related Commands	show system profile	
Notes		

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
Role	admin
Example	switch (config) # show system mac 00:02:C9:5E:AF:18 switch (config) #
Related Commands	N/A
Notes	This command is only available in Ethernet system profile

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
Role	admin
Example	switch (config) # show system profile eth-single-switch switch (config) #
Related Commands	system profile
Notes	

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
Role	admin
Example	switch (config) # show system type SX1036 switch (config) #
Related Commands	
Notes	

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
Role	admin
Example	<pre> switch (config) # show temperature ===== Module Component Reg CurTemp Status (Celsius) ===== MGMT BOARD_MONITOR T1 25.00 OK MGMT CPU_BOARD_MONITOR T1 26.00 OK MGMT CPU_BOARD_MONITOR T2 41.00 OK MGMT QSFP_TEMP1 T1 23.00 OK MGMT QSFP_TEMP2 T1 22.50 OK MGMT QSFP_TEMP3 T1 23.00 OK MGMT SX T1 37.00 OK switch (config) # </pre>
Related Commands	N/A
Notes	

show uboot

show uboot

Displays u-boot version.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.5006
	3.4.1110 Updated Example
Role	admin
Example	switch (config) # show uboot UBOOT version : U-Boot 2009.01 SX_PPC_M460EX 3.2.0330-82 ppc (Dec 20 2012 - 17:53:54) switch (config) #
Related Commands	N/A
Notes	This command is available only on PPC based systems

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
Role	admin
Example	<pre> switch (config) # show version Product name: MLNX-OS Product release: 3.1.0000 Build ID: #1-dev Build date: 2012-02-26 08:47:51 Target arch: ppc Target hw: m460ex Built by: root@r-fit16 Uptime: 1d 3h 32m 24.656s Product model: ppc Host ID: 0002c911a15e System memory: 110 MB used / 1917 MB free / 2027 MB total Swap: 0 MB used / 0 MB free / 0 MB total Number of CPUs: 1 CPU load averages: 0.18 / 0.19 / 0.16 switch (config) # </pre>
Related Commands	N/A
Notes	

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
Role	admin
Example	<pre>switch (config) # show version concise PPC_M460EX 3.4.2000 2015-05-06 20:26:41 ppc switch (config) #</pre>
Related Commands	N/A
Notes	

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
Role	admin
Example	<pre>switch (config) # show voltage ===== Module Power Meter Reg Expected Actual Status High Low Voltage Voltage Voltage ===== 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 switch (config) #</pre>
Related Commands	N/A
Notes	

4.17.8.2 Chassis High Availability

chassis ha bip

chassis ha bip <board IP address>

Configures Chassis Board IP (BIP).

Syntax Description	board IP address	Sets the chassis virtual IP address.
Default	0.0.0.0	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # chassis ha bip 192.168.10.100 switch (config) #</pre>	
Related Commands	show chassis ha	
Notes	This command is applicable only for director switch systems.	

chassis ha

chassis ha reset other

Performs a reset to the other management card in the chassis.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # chassis ha reset other switch (config) #
Related Commands	show chassis ha
Notes	This command is applicable only for director switch systems.

chassis ha power enable other

chassis ha power enable other
no chassis ha power enable other

Enables the other management card in the chassis.
The no form of the command disables the other management card in the chassis.

Syntax Description	N/A
Default	The other management card is enabled.
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # chassis ha power enable other switch (config) #
Related Commands	show chassis ha
Notes	This command is applicable only for director switch systems.

show chassis ha

show chassis ha

Displays Chassis HA parameters and status.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # show chassis ha 2-node HA state: Box management IP: 172.30.1.200/16 interface: mgmt0 local role: master local slot: 1 other state: ready reset count: 0 switch (config) #</pre>
Related Commands	chassis ha
Notes	This command is applicable only for director switch systems.

show chassis ha

show chassis ha

Displays Chassis HA parameters and status.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # show chassis ha 2-node HA state: Box management IP: 172.30.1.200/16 interface: mgmt0 local role: master local slot: 1 other state: ready reset count: 0 switch (config) #</pre>
Related Commands	chassis ha
Notes	This command is applicable only for director switch systems.

4.18 Network Management Interfaces

4.18.1 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.

MLNX-OS supports:

- SNMP versions v1, v2c and v3
- SNMP trap notifications
- Standard MIBs
- Mellanox private MIBs

4.18.1.1 Standard MIBs

Table 33 - Standard MIBs – Textual Conventions and Conformance MIBs

MIB	Standard	Comments
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		



Starting from version 3.4.1600, IB interfaces in interfaces tables (i.e. ifTable, ifxTable) have changed from SX<if>/<port> to IB/port.

Table 34 - Standard MIBs – Chassis and Switch

MIB	Standard	Comments
RFC1213-MIB	RFC 1213	
IF-MIB	RFC 2863	ifXTable only supported.
ENTITY-MIB	RFC 4133	
ENTITY-SENSOR-MIB	RFC 3433	Fan and temperature sensors

Table 34 - Standard MIBs – Chassis and Switch

MIB	Standard	Comments
ENTITY-STATE-MIB	RFC 4268	Fan and temperature states
Bridge MIB	RFC 4188	dot1dTpFdbGroup and dot1dStaticGroup are not supported in this MIB, it is supported as a part of Q-Bridge-MIB. This MIB is not relevant to InfiniBand.
Q-Bridge MIB	RFC 4363	The following SNMP groups are not supported: • qBridgeVlanStatisticsGroup, • qBridgeVlanStatisticsOverflowGroup , • qBridgeVlanHCStatisticsGroup, • qBridgeLearningConstraintsGroup. The following SNMP tables are not supported: • dot1qTpGroupTable (dynamic MC MAC addresses) • dot1qForwardAllTable (GMRP) • dot1qForwardUnregisteredTable (GMRP) • dot1qVlanCurrentTable (GVRP) This MIB is not relevant to InfiniBand.
RSTP-MIB	RFC 4318	This MIB is not relevant to InfiniBand.
LLDP-MIB	802.1AB-2005	This MIB is not relevant to InfiniBand.
ENTITY-SENSOR-MIB	RFC 3433	• Port module transmit/receiver power sensors (for 1U systems only) • Fan, temperature sensors
BGP4-MIB	RFC 4273	Only supports the following tables: • bgpLocalAs • bgpPeerLocalAddr • bgpPeerState • bgpIdentifier This MIB is not relevant to InfiniBand.
OSPF-MIB	RFC 4750	This MIB is not relevant to InfiniBand.

4.18.1.2 Private MIB

Table 35 - Private MIBs Supported

MIB	Description
MELLANOX-SMI-MIB	Mellanox 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
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 private MIBs can be downloaded from the [Mellanox Support](#) webpage.

4.18.1.3 Mellanox Private Traps

The following private traps are supported by MLNX-OS®.

Table 36 - SNMP Traps

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 Mellanox support.
cpuUtilHigh	N/A
procUnexpectedExit	Generate SysDump and contact Mellanox support.
diskSpaceLow	Clean images and sysDump files using the commands “image delete” and “file debug-dump delete”.
systemHealthStatus	Refer to Health Status table.

Table 36 - SNMP Traps

Trap	Action Required
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 [Table 28, “Supported Event Notifications and MIB Mapping,”](#) on page 351.

4.18.1.4 Configuring SNMP

➤ *To set up the SNMP:*

- Step 1.** Activate the SNMP server on the MLNX-OS switch (in configure mode) using the following commands:



Community strings are case sensitive.



Director switches require SNMP timeout configuration on the agent of 60 seconds.

```
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
```

4.18.1.5 Configuring an SNMPv3 User

➤ *To configure SNMPv3 user:*

- Step 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
- priv type – des or aes-128

Step 2. Enter authentication password and its confirmation.

Step 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: *****
switch (config) #
```

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
```

4.18.1.6 Configuring an SNMP Notification

➤ *To set up the SNMP Notification (traps or informs):*

Step 1. Make sure SNMP and SNMP notification are enable. Run:

```
switch (config) # snmp-server enable
switch (config) # snmp-server enable notify
switch (config) #
```

Step 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
switch (config) #
```

Step 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)

switch (config) #
```

Step 4. Configure the desired event to be sent via SNMP. Run:

```
switch (config) # snmp-server notify event interface-up
switch (config) #
```



This particular event is used as an example only.

Step 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 triggerred a topology change
unexpected-shutdown: Unexpected system shutdown
switch (config) #
```



To print event notifications to the terminal (SSH or CONSOLE) refer to [Section 4.8.1, “Monitor,”](#) on page 306.

4.18.1.7 SNMP SET Operations

MLNX-OS allows the user to use SET operations via SNMP interface. This is needed to configure a user/community supporting SET operations.

4.18.1.7.1 Enabling SNMP SET

➤ **To allow SNMP SET operations using SNMPv1/v2:**

Step 1. Enable SNMP communities. Run:

```
switch (config) # snmp-server enable communities
```

Step 2. Configure a read-write community. Run:

```
switch (config) # snmp-server community my-community-name rw
```

- Step 3.** Make sure SNMP communities are enabled (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
switch (config) # show snmp
No Listen Interfaces.
```

- Step 4.** Configure this RW community in your MIB browser.

➤ **To allow SNMP SET operations using SNMPv3:**

- Step 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](#)” on [page 579](#) for more information.

- Step 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
```

MLNX-OS supports the OIDs for SET operation listed in Table 37 which are expanded upon in the following subsections.

Table 37 - Supported SET OIDs

MIB Name	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

Table 37 - Supported SET OIDs

MIB Name	OID Name	OID
MELLANOX-CONFIG-DB	mellanoxConfigDBCmdExecute	1.3.6.1.4.1.33049.12.1.1.2.3.0
	mellanoxConfigDBCmdFilename	1.3.6.1.4.1.33049.12.1.1.2.2.0
	mellanoxConfigDBCmdStatus	1.3.6.1.4.1.33049.12.1.1.2.4.0
	mellanoxConfigDBCmdStatusString	1.3.6.1.4.1.33049.12.1.1.2.5.0
	mellanoxConfigDBCmdUri	1.3.6.1.4.1.33049.12.1.1.2.1.0
MELLANOX-POWER-CYCLE	mellanoxPowerCycleCmdExecute	1.3.6.1.4.1.33049.10.1.1.2.1.0
	mellanoxPowerCycleCmdStatus	1.3.6.1.4.1.33049.10.1.1.2.2.0
	mellanoxPowerCycleCmdStatusString	1.3.6.1.4.1.33049.10.1.1.2.3.0
MELLANOX-SW-UPDATE	mellanoxSWUpdateCmdSetNext	1.3.6.1.4.1.33049.11.1.1.2.1.0
	mellanoxSWUpdateCmdUri	1.3.6.1.4.1.33049.11.1.1.2.2.0
	mellanoxSWUpdateCmdExecute	1.3.6.1.4.1.33049.11.1.1.2.3.0
	mellanoxSWUpdateCmdStatus	1.3.6.1.4.1.33049.11.1.1.2.4.0
	mellanoxSWUpdateCmdStatusString	1.3.6.1.4.1.33049.11.1.1.2.5.0
	mellanoxSWActivePartition	1.3.6.1.4.1.33049.11.1.1.3.0.0
	mellanoxSWNextBootPartition	1.3.6.1.4.1.33049.11.1.1.4.0.0

4.18.1.7.2 Sending a Test Trap SET Request

MLNX-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 [Section 4.18.1.7.1, “Enabling SNMP SET,”](#) on page 553.
2. Configure host to which to send SNMP notifications.
3. Set a trap receiver in the MIB browser.

➤ **To send a test trap:**

Step 1. Send a SET request to the switch IP with the OID 1.3.6.1.4.1.33049.2.1.1.6.0.

Step 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).

4.18.1.7.3 Setting Hostname with SNMP

Mellanox 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”](#) on page 180 for more information.

4.18.1.7.4 Power Cycle with SNMP

Mellanox supports power cycling its systems using an SNMP SET request as described in MEL-LANOX-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`)
- Reload slave – reloads the slave management on dual management systems (must be executed from the master management module)



On dual management systems it is advised to connect via the BIP to make sure commands are executed from the master management.

4.18.1.7.5 Changing Configuration with SNMP

Mellanox 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 `mellanoxConfigD-`

BCmdUri 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 mellanoxConfigD-BCmdUri 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 or an existing configuration file on the switch to the URI in the mellanoxConfigD-BCmdUri 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

4.18.1.7.6 Upgrading MLNX-OS Software with SNMP

Mellanox supports upgrading MLNX-OS 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” on page 262](#).

- 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%
- 101 – success
- 200 – failure

4.18.1.8 IF-MIB and Interface Information

MLNX-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. The user has the ability 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).

4.18.2 JSON API



The JSON API is available on x86 based systems only.

JavaScript Object Notation (JSON) is a machine-to-machine data-interchange format which is supported in MLNX-OS® CLI.

The JSON API allows executing CLI commands and receiving outputs in JSON format which can be easily parsed by the calling software.

4.18.2.1 JSON Request Over HTTP/HTTPS

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.



HTTP access to the web-based management console needs to be enabled using the command “web http enable” to allow POST requests.

4.18.2.1.1 Authentication

The HTTP client must first be authenticated by sending a POST request to the following URL:

<http://<switch-ip-address>/admin/launch?script=rh&template=login&action=login>.

The POST request content should contain the following variables:

- f_user_id
- f_password

4.18.2.1.2 Sending the Request

After authenticating the HTTP client, the user needs to construct a POST request containing the CLI command they wish to execute as the payload of the request (see [Section 4.18.2.1.3, “JSON Request Format,” on page 559](#) for the CLI request format). Using the HTTP client, the POST request should be sent to the following URL: <http://<switch-ip-address>/admin/launch?script=json>.

After the request is handled in the system the HTTP client receives a reply with an indication of the request execution result. If there is data resulting from the request, it is returned as part of the reply.

See for the reply format in [Section 4.18.2.1.4, “JSON Response Format,” on page 559](#).

JSON POST requests may also be sent using the WebUI. For more information, please refer to [Section 4.18.2.2, “JSON Request Using WebUI,” on page 562](#).

4.18.2.1.3 JSON Request Format

The request format is a JSON object consisting of a single name-value pair where the name is the string “cmd” and the value is the CLI command the user executed as a string also.

```
{
  "cmd": "<CLI command to execute>"
}
```

Only one query/request command is supported at a time.

See [Section 4.18.2.1.6, “JSON Examples,” on page 560](#) for usage example.

4.18.2.1.4 JSON Response Format

The HTTP POST response format structure is a JSON object consisting of 3 name-value pairs as follows:

```
{
  "status" = "<OK|ERROR>",
  "data" = {The information requested}
  "status_message" = "<Information on the status received>",
}
```

- 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”.

4.18.2.1.5 Supported 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.

Not all CLI show commands are currently supported by the JSON API. Unsupported commands return an error indication.

Support for all show commands will be completed in future MLNX-OS releases.

For a list of currently supported “show” commands, please refer to [Appendix C, “Show Commands Supported by JSON API,” on page 1646](#).

4.18.2.1.6 JSON Examples

The following examples use cURL (common tool in Linux systems) to send HTTP POST requests to the system.

Authentication Example

Before sending JSON HTTP request, the user must first authenticate. Run the following from your server’s shell to create a login session ID in the file: /tmp/cookie.

```
curl -c /tmp/cookie -d "f_user_id=admin&f_password=admin"
"http://10.10.10.10/admin/launch?script=rh&template=login&action=login"
```

Query Request Example

This example sends a request to query the XML gateway setting on the switch system.

Step 1. Prepare your request.

The request payload must look like this:

```
{"cmd": "show xml-gw"}
```

Send the POST request using cURL with the request payload as follows:

```
echo '{"cmd": "show xml-gw"}' | curl -b /tmp/cookie -X POST -d @-
"http://10.10.10.10/admin/launch?script=json"
```


Step 2. Get the reply.

After sending the request, when the system finishes processing, a reply similar to the following is received:

```
{
  "status": "OK",
  "data": {
    "XML Gateway enabled": "yes"
  },
  "status_message": ""
}
```

You may also prepare a file containing the JSON request as follows:

Step 1. Prepare the JSON request.

The file used in this example is “req_example.json”:

```
{
  "cmd": "show xml-gw"
}
```

Step 2. Use it as the POST request payload like this. Run:

```
curl -b /tmp/cookie -X POST -d @req_example.json http://10.10.10.10/admin/
launch?script=json
```

Set Request Example

This example sends a request to enable SNMP server.

Step 1. Prepare your request

The request payload should look like this:

```
{"cmd": "snmp-server enable"}
```

Send the POST request using cURL with the request payload as follows:

```
echo '{"cmd": "snmp-server enable"}' | curl -b /tmp/cookie -X POST -d @-
"http://10.10.10.10/admin/launch?script=json"
```

Step 2. Get the reply.

After sending the request, when the system finishes processing, a reply similar to the following is received:

```
{
  "status": "OK",
  "data": "",
  "status_message": ""
}
```



Set commands normally do not return any data or output. In instances where a set command returns an output, it is displayed in the "status_message" field.

Error Response Example

If the request fails, the status value returned is “ERROR” and, if applicable, the “status_msg” value will contain a description of the problem.

For example, sending a request without authentication results in the following ERROR:

```
echo '{"cmd":"show system profile"}' | curl -b /tmp/cookie -X POST -d @-  
"http://10.10.10.10/admin/launch?script=json"  
{  
  'status': 'ERROR'  
  'data': ''  
  'status_message': 'Request not authenticated'  
}
```

4.18.2.2 JSON Request Using WebUI

The MLNX-OS® WebUI also allows users to send JSON HTTP POST requests.


Login to the WebUI → Go to the “Setup” tab → Select “JSON API” from the left side menu.



This section is displayed only if JSON API is enabled using the command “json-gw enable”.

In the “cmd” field type the CLI command to execute then press “Submit”. The JSON response is then shown in the “JSON Response” box below.

Figure 18: JSON API WebUI Example



Mellanox MLNX-OS Management Console
Host: . User: admin Logout

Standalone

Setup System Security Ports **JSON API** IB SM Mgmt Fabric Inspector ETH Mgmt IP Route Gateway Save

JSON API Product Documents

Interfaces
HA
Routing
Hostname
DNS
Login/Logout Messages
Address Resolution
IPSec
Neighbors
Virtualization
Virtual Switch Mgmt
Web
SNMP
Email Alerts
XML gateway
JSON API
Logs
Configurations
Date and Time
NTP
Licensing

JSON Configuration

Enable JSON API ☒

Apply Cancel

Execute JSON Command

Enter a CLI command to be executed:

```
{
  "cmd": "show xml-gw"
}
```

Submit Cancel

JSON Response:

```
{
  "status": "OK",
  "data": {
    "XML Gateway enabled": "yes"
  },
  "status_message": ""
}
```

4.18.3 XML API

MLNX-OS XML API is documented in the *MLNX-OS® XML API Reference Guide*.

4.18.4 Commands

4.18.4.1 SNMP Commands

The commands in this section are used to manage the SNMP server.

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	In seconds. Range: 20-500.
Default	Enabled. Interval: 60 secs	
Configuration Mode	Config	
History	3.2.3000	
	3.4.1100	Added time parameter and updated notes
Role	admin	
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.	

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	
Role	admin	
Example	<pre> switch(config) # snmp-server community private rw switch (config) # show snmp SNMP enabled: yes SNMP port: 161 System contact: System location: Read-only community: public Read-write community: private Interface listen enabled: yes No Listen Interfaces. Traps enabled: yes Default trap community: public Default trap port: 162 No trap sinks configured. switch(config) # </pre>	
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	

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	
Role	admin	
Example	<pre> switch (config) # snmp-server contact my-name switch (config) # show snmp SNMP enabled: yes SNMP port: 161 System contact: my-name System location: Read-only community: public Read-write community: private Interface listen enabled: yes No Listen Interfaces. Traps enabled: yes Default trap community: public Default trap port: 162 No trap sinks configured. switch (config) # </pre>	
Related Commands	show snmp	
Notes		

snmp-server enable

snmp-server enable
no snmp-server enable

Enables SNMP-related functionality (SNMP engine, and traps)
The no form of the command disables the SNMP server.

Syntax Description	N/A
Default	SNMP is enabled by default
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # snmp-server enable
Related Commands	show snmp
Notes	

snmp-server enable communities

snmp-server enable communities
no snmp-server enable communities

Enables community-based authentication.
The no form of the command disables community-based authentication.

Syntax Description	N/A
Default	SNMP server communities are enabled by default
Configuration Mode	Config
History	N/A
Role	admin
Example	switch (config) # snmp-server enable communities
Related Commands	show snmp
Notes	

snmp-server enable multi-communities

snmp-server enable multi-communities
no snmp-server enable multi-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	N/A
Role	admin
Example	switch (config) # snmp-server enable multi-communities
Related Commands	show snmp
Notes	

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	N/A
Role	admin
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.

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
Role	admin
Example	switch (config) # snmp-server enable set-permission MELLANOX-SW-UPDATE
Related Commands	show snmp set-permission
Notes	

snmp-server host

snmp-server host <IP address> {disable | {traps | informs} [<community> | <port> | version <snmp version>]}

no snmp-server host <IPv4 or IPv6 address> {disable | {traps| informs} [<community> | <port>]}

Configures hosts to which to send SNMP traps.

The no form of the commands removes a host from which SNMP traps should be sent.

Syntax Description	IP address	IPv4 or IPv6 address.
	disable	Temporarily disables sending of traps to this host.
	community	Specifies trap community string.
	port	Overrides default UDP port for this trap sink.
	snmp version	Specifies the SNMP version of traps to send to this host.
Default	No hosts are configured Default community is “public” Default UDP port is 162 Default SNMP version is 2c	
Configuration Mode	Config	
History	3.1.0000	First version
	3.2.1050	Add inform option
Role	admin	

Example

```
switch (config) # snmp-server host 10.10.10.10 traps version 1
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.

Traps enabled:          yes
Default trap community: public
Default trap port:      162

Trap sinks:
    10.10.10.10
        Enabled: yes
        Type: traps version 1
        Port: 162 (default)
        Community: public (default)
switch (config) #
```

Related Commands

```
show snmp
snmp-server enable
```

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 “snmp-server enable” command.

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	
Role	admin	
Example	<pre> switch (config) # snmp listen enable switch (config) # show snmp SNMP enabled: yes SNMP port: 161 System contact: System location: Read-only community: public Read-write community: private Interface listen enabled: yes No Listen Interfaces. Traps enabled: yes Default trap community: public Default trap port: 162 Trap sinks: 10.10.10.10 Enabled: yes Type: traps version 1 Port: 3 Community: public (default) switch (config) # </pre>	
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.	

snmp-server location

snmp-server location <system location>
no snmp-server location

Sets a value for the sysLocation variable in MIB-II.
The no form of the command clears the contents of the sysLocation variable.

Syntax Description	system location	String.
Default	""	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # snmp-server location lab switch (config) # show snmp SNMP enabled: yes SNMP port: 161 System contact: my-name System location: lab Read-only community: public Read-write community: private Interface listen enabled: yes No Listen Interfaces. Traps enabled: yes Default trap community: public Default trap port: 162 No trap sinks configured. switch (config) #</pre>	
Related Commands	show snmp	
Notes		

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	Community: public All informs and traps are enabled Port: 162	
Configuration Mode	Config	
History	3.1.0000	First version
	3.2.1050	Changed traps to notify
Role	admin	
Example	<pre>switch (config) # snmp-server community public switch (config) # show snmp SNMP enabled: yes SNMP port: 1000 System contact: my-name System location: lab Read-only community: public Read-write community: private Interface listen enabled: yes No Listen Interfaces. Traps enabled: yes Default trap community: public Default trap port: 162 No trap sinks configured. switch (config) #</pre>	

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
-

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	
Role	admin	
Example	<pre>switch (config) # snmp-server port 1000 switch (config) # show snmp SNMP enabled: yes SNMP port: 1000 System contact: my-name System location: lab Read-only community: public Read-write community: private Interface listen enabled: yes No Listen Interfaces. Traps enabled: yes Default trap community: public Default trap port: 162 No trap sinks configured. switch (config) #</pre>	
Related Commands	show snmp	
Notes		

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 SNMP v3 users
	auth	Configures SNMP v3 security parameters, specifying passwords in plaintext on the command line (note: passwords are always stored encrypted)
	capability	Sets capability level for SET requests
	enable	Enables SNMP v3 access for this user
	encrypted	Configures SNMP v3 security parameters, specifying passwords in encrypted form
	prompt	Configures SNMP v3 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
Default	No SNMP v3 users defined	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # snmp-server user admin v3 enable switch (config) # show snmp user User name: admin Enabled overall: yes Authentication type: sha Privacy type: aes-128 Authentication password: (NOT SET; user disabled) Privacy password: (NOT SET; user disabled) SET access: Enabled: yes Capability level: admin switch (config) #</pre>	

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.
-

show snmp

show snmp [auto-refresh | engineID | events | host | user]

Displays SNMP-server configuration and status.

Syntax Description	auto-refresh	SNMP refreshed mechanism status.
	engineID	SNMP Engine ID.
	events	SNMP events.
	host	List of notification sinks.
	user	SNMP users.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>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 switch (config) #</pre>	
Related Commands	show snmp	
Notes		

show snmp auto-refresh

show snmp auto-refresh

Displays SNMPD refresh mechanism status.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre> switch (config) # show snmp auto-refresh ===== SNMP auto refresh ===== Auto-refresh enabled: yes Refresh interval (sec): 60 ===== Auto-Refreshed tables ===== entPhysicalTable ifTable ifXTable </pre>
Related Commands	snmp-server auto-refresh
Notes	

show snmp set-permission

show snmp set-permission

Displays SNMP SET permission settings.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.6.3004
Role	admin
Example	<pre>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</pre>
Related Commands	snmp-server enable set-permission
Notes	

show interfaces ib internal notification

show interfaces ib internal notification

Displays information about internal links notification.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.4318
	3.4.3000 Updated output
Role	admin
Example	<pre>switch (config) # show interfaces ib internal notification ===== Internal links information ===== State change enabled : yes Speed mismatch enabled : yes Periodic notifications : 6 (hours)</pre>
Related Commands	interfaces ib internal notification
Notes	

4.18.4.2 XML API Commands

xml-gw enable

xml-gw enable
no xml-gw enable

Enables the XML gateway.
The no form of the command disables the XML gateway.

Syntax Description	N/A
Default	XML Gateway is enabled
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # xml-gw enable switch (config) # show xml-gw XML Gateway enabled: yes switch (config) #
Related Commands	show xml-gw
Notes	

show xml-gw

show xml-gw

Displays the XML gateway setting.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.1.0000
Role	admin
Example	switch (config) # show xml-gw XML Gateway enabled: yes switch (config) #
Related Commands	xml-gw enable
Notes	

4.18.4.3 JSON API Commands

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
Role	admin
Example	switch (config) # json-gw enable
Related Commands	show json-gw
Notes	This command is available on x86 switch systems only.

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
Role	admin
Example	switch (config) # show json-gw JSON Gateway enabled: yes
Related Commands	json-gw enable
Notes	This command is available on x86 switch systems only.

4.19 Puppet Agent

Puppet is a software that allows network administrators to automate repetitive tasks. MLNX-OS includes a built-in agent for the open-source “Puppet” configuration change management system. The Puppet agent enables configuring Mellanox switches in accordance with the standard “puppet-netdev-stdlib” type library and with the “Mellanox-netdev-stdlib-mlnxos” and “Mellanox-netdev-ospf-stdlib” type libraries provided by Mellanox Technologies to the Puppet community.

For more information, please refer to the CLI commands, to the NetDev documentation at <https://github.com/puppetlabs/puppet-netdev-stdlib> and to Mellanox’s Puppet modules GitHub page at <https://github.com/Mellanox>.

4.19.1 Setting the Puppet Server

➤ *To set the puppet server:*

Step 1. Define the Puppet server (the name has to be a DNS and not IP). Run:

```
switch (config) # puppet-agent master-hostname <please_type_your_hostname_DNS_here>
switch (config) #
```

Step 2. Enable the Puppet agent. Run:

```
switch (config) # puppet-agent enable
switch (config) #
```

Step 3. (Optional) Verify there are no errors in the Puppet agent log. Run:

```
switch (config) # show puppet-agent log continuous
switch (config) #
```

4.19.2 Accepting the Switch Request



This is to be performed on the first run only.

➤ *To accept the switch’s request:*

Option 1 – using Puppet CLI commands:

Step 1. Ensure the certificate request. Run:

```
# puppet cert list
"<switch>"
(F4:B4:20:3B:2B:11:76:37:14:34:D0:D1:03:ED:3D:B5)
```

Step 2. Sign the certificate request if the cert_name parameter (e.g. switch1.domain) is in the list. Run:

```
# puppet cert sign <full_domain_name>
```

Step 3. Verify the request is removed from the Puppet certification list. Run:

```
# puppet cert list
```

Option 2 – accept certificate requests in the puppet server console:

- Step 1.** Go to the “nodes requests” page (the button is at the top right), and wait for a certificate request for the switch and then accept it.

Figure 19: Accepting an Agent Request through the Console



4.19.3 Installing Modules on the Puppet Server

Mellanox uses netdev-stdlib types and provides a package of Mellanox providers for those types which have to be installed at the Puppet server prior to the first Puppet configuration run (before configuring resources on the Mellanox switch).

To install those modules, run the following commands in the Puppet server:

```
# puppet module install netdevops-netdev_stdlib
# puppet module install mellanox-netdev_ospf_stdlib
# puppet module install mellanox-netdev_stdlib_mlnxos
```



In case of an already installed module, please use the command “puppet module upgrade <module_name>” or “puppet module install <module_name> -force” instead of “puppet module install <module_name>” to reinstall the modules.

For more information please refer to the Network Automation Tools document or Puppet category in the Mellanox community site at: <http://community.mellanox.com/community/support/solutions>.

4.19.4 Writing Configuration Classes

➤ **To write configuration classes:**

- Step 1.** Assigning Configuration Classes to a Node

Configuration files can be written and changed in the puppet server machine in the directory “/etc/puppetlabs/puppet/manifests/” (or “/etc/puppet/manifests” in case of an open source puppet server).

The file “/etc/puppetlabs/puppet/manifests/site.pp” is the main file for Puppet-classes-to-nodes association. To associate a configuration to a Puppet agent node, just append association lines as below:

```
import "netdev_vlan_example"
import "netdev_l2_vlan_example"
import "netdev_lag_example"
node 'switch-6375dc.mtr.labs.mlnx'{

    netdev_device { $hostname: }

    include vlan_example # Asserts a class vlan_example in one of the files
    include l2_interface_example

    include lag_example

}
```



If you have a puppet console, you may assign classes of configuration in the following way:

- Add the relevant classes (using the console add class button on the “nodes” page).
- Assign the classes to the relevant nodes/groups in the puppet server console (in the console node/group page -> edit -> Classes).

Step 2. Update VLAN

Manifest example (located in “/etc/puppetlabs/puppet/manifests/netdev_vlan_example.pp”).

```
class vlan_example{

    $vlans = {
        'Vlan244' => {vlan_id => 244, ensure => present},
        'Vlan245' => {vlan_id => 245, ensure => present},
    }

    create_resources( netdev_vlan, $vlans )
}
```

Step 3. Update Layer 2 Interface.

Manifest example (located in “/etc/puppetlabs/puppet/manifests/netdev_l2_interface_example.pp”)

```
class vlans_ensure_example{

    $vlans = {
        'Vlan347' => {vlan_id => 347, ensure => present},
        'Vlan348' => {vlan_id => 348, ensure => present},
        'Vlan349' => {vlan_id => 349, ensure => present},
    }

    create_resources( netdev_vlan, $vlans )
}

class l2_interface_example{

    include vlans_ensure_example #class to Ensure VLANs before assigning

    $l2_interfaces = {
        'ethernet 1/3' => {ensure => absent, vlan_tagging => disable}, #default
        'ethernet 1/4' => {ensure => present, vlan_tagging => enable,
        tagged_vlans => [Vlan348,Vlan347], untagged_vlan => Vlan349} #hybrid
    }

    create_resources( netdev_l2_interface, $l2_interfaces )
}
```

Step 4. Update LAG.

Manifest example (located in “/etc/puppetlabs/puppet/manifests/netdev_lag_example.pp”)

```
class lag_example{

    $lags = {
        'port-channel 101' => {ensure => present,
        links => ['ethernet 1/12', 'ethernet 1/13'], lacp => active},
        'port-channel 102' => {ensure => present,
        links => ['ethernet 1/6','ethernet 1/5'], lacp => disabled},
    }

    create_resources( netdev_lag, $lags )
}
```



You may add classes to ensure that all assigned links are with the same layer 1 and layer 2 configurations (similarly to the way we did in update l2_interface section with vlans_ensure_example class).

4.19.5 Supported Configuration Capabilities

4.19.5.1 Ethernet, InfiniBand, and Port-Channel Interface Capabilities

Table 38 - Ethernet and Port-Channel Interface Capabilities

Field	Description	Values	Example
ensure	Sets the given values or restores the interface to default	absent, present	ensure => present
speed	Sets the speed of the interface.	auto* 10m 100m 1g 10g 40g 56g	speed => 1g
admin	Disables/enables interface admin state.	up, down	admin => up
mtu	Configures the maximum transmission unit frame size for the interface.	Ethernet: 1518-9216	mtu => 1520
description	Sets the Ethernet and LAG description.	Text	description => "changed_by_puppet"

4.19.5.2 VLAN Capabilities

Table 39 - VLAN Capabilities

Field	Description	Values	Example
ensure	Creates or destroys the VLAN given as a resource ID	absent, present	ensure => present
vlan_id	The VLAN ID	1-4094 (integer)	vlan_id => 245

4.19.5.3 Layer 2 Ethernet Interface Capabilities

Table 40 - L2 Ethernet and Port-Channel Interface Capabilities

Field	Description	Values	Example
ensure	Sets the given values or restores the Layer 2 interface to default.	absent, present	ensure => present
vlan_tagging	VLAN tagging mode	enable,disable	vlan_tagging => enable
tagged_vlans	List of tagged (trunked) VLANs	2-4994 (range)	tagged_vlans => [Vlan348,Vlan347]
untagged_vlan	Untag (access) VLAN	<VLAN name>	untagged_vlan => Vlan349

4.19.5.4 LAG (Port-Channel) Capabilities

Table 41 - LAG Capabilities

Field	Description	Values	Example
ensure	creates or destroys the port-channel given as a resource ID	absent, present	ensure => present
lacp	The LACP mode of the LAG	passive active on	lacp => on
links	List of ports assigned to the LAG	List of link names	links => ['ethernet 1/6','ethernet 1/5']

4.19.5.5 Layer 3 Interface Capabilities

Table 42 - L3 Interface Capabilities

Field	Description	Values	Example
ensure	Creates or destroys the interface VLAN specified in the resource ID.	present, absent	ensure => present
ipaddress	Sets IP address on the Layer 3 interface (requires netmask).	A valid IP address	ipaddress => '192.168.4.2'
netmask	Sets netmask for the IP address.	A valid netmask (of the form X.1X2.X3.X4), which creates a valid combination with the given IP address	netmask => '255.255.255.0'
method	Configures the method of the L3 interface (currently supports only static method).	static	method => static

4.19.5.6 OSPF Interface Capabilities

Table 43 - OSPF Interface Capabilities

Field	Description	Values	Example
ensure	Creates or destroys the OSPF interface of the associated interface of the VLAN specified in the resource ID	present, absent	ensure => present
area_id	The associated area ID	Integer representing an IP	area_id => '7200'

Table 43 - OSPF Interface Capabilities

Field	Description	Values	Example
Type	The network type	broadcast, point_to_point	type => 'point_to_point'

4.19.5.7 OSPF Area Capabilities

Table 44 - OSPF Area Capabilities

Field	Description	Values	Example
ensure	Creates or destroys the OSPF area specified in the resource ID	present, absent	ensure => present
router_id	The OSPF area associated router ID (currently supports only default router)	default	router_id => 'default'
ospf_area_mode	The OSPF area mode	normal, stub, nssa	ospf_area_mode => 'stub'
subnets	A list of associated subnets	List of subnets	["192.168.4.0/24", "192.168.5.0/24"]

4.19.5.8 Router OSPF Capabilities

Table 45 - Router OSPF Capabilities

Field	Description	Values	Example
ensure	Enables/disables the router ID specified in the resource ID	present, absent	ensure => present

4.19.5.9 Protocol LLDP, SNMP, IP Routing and Spanning Tree Capabilities

Table 46 - Protocol Enable/Disable Capabilities

Field	Description	Values	Example
ensure	Enables/disables the protocol specified in the resource ID	present, absent	ensure => present

4.19.5.10 Fetched Image Capabilities

Table 47 - Fetched Image Capabilities

Field	Description	Values	Example
ensure	Enables/disables the protocol specified in the resource ID	present, absent	ensure => present
protocol	Specifies the protocol for fetch method	http, https, ftp, tftp, scp, sftp	protocol => scp
host	The host where the file-name located	DNS/IP	host => my_DNS

Table 47 - Fetched Image Capabilities

Field	Description	Values	Example
user	The username for fetching the image	Username	user => my_username
password	The password for fetching the image	Password	password => my_pass
location	The location of the file name in the host file system	Directory full path	location => '/tmp'
force_delete	Remove all the images or only the ones which are not installed on any partition, before fetching	yes, no	force_delete => no

4.19.5.11 Installed Image Capabilities

Table 48 - Installed Image Capabilities

Field	Description	Values	Example
ensure	Specifies if the image version given in as resource ID is ensured to be installed or not	present, absent	ensure => present
is_next_boot	Ensures that the installed image is the next boot partition	yes, no	is_next_boot => yes
configuration_write	Writes configurations to database.	yes, no	configuration_write => yes
force_reload	Reload if image is in other partition.	yes, no	force_reload => no

4.19.6 Supported Resources for Each Type

Table 49 - Fetched Image Capabilities

Resource Type	Puppet Type Name	Supported Resource IDS	Example
Network device	netdev_device	\$hostname	netdev_device { \$hostname: }
Layer 1 interface	netdev_interface	'ethernet <#ID>', 'port-channel <#id>', 'ib <#ID>'	netdev_interface { 'ethernet 1/3': ensure => absent }

Table 49 - Fetched Image Capabilities

Resource Type	Puppet Type Name	Supported Resource IDS	Example
Layer 2 interface	netdev_l2_interface	'ethernet <#ID>', 'port-channel <#id>'	netdev_l2_interface { 'ethernet 1/3': ensure => absent }
VLAN	netdev_vlan	VLAN name string	netdev_vlan { 'Vlan244': vlan_id => 244, ensure => present }
LAG	netdev_lag	'port-channel <#id>'	netdev_lag { 'port-channel 101': ensure => present }
Layer 3 interface	netdev_l3_interface	'vlan <#ID>'	netdev_l3_interface { 'vlan 4': ipaddress => '192.168.4.2', netmask => '255.255.255.0' }
OSPF interface	netdev_ospf_interface	'vlan <#ID>'	netdev_ospf_interface { 'vlan 4': ensure => present, area_id => '10' }
OSPF area	netdev_ospf_area	Valid area ID (representing an IP)	netdev_ospf_area { '10': ensure => present, ospf_area_mode => 'stub' }
OSPF router	netdev_router_ospf	Currently only supports 'default'	netdev_router_ospf { 'default': ensure => present }
Protocol	mlnx_protocol	ip_routing, lldp, snmp, spanning_tree	mlnx_protocol { 'ip_routing': ensure => present }
Fetched image	mlnx_fetched_img	The image file name	mlnx_fetched_image { 'image-PPC_M460EX-3.3.4300.img': ensure => present }
Installed image	mlnx_installed_img	The image version name	mlnx_installed_img { '3.3.4300': ensure => present }

4.19.7 Troubleshooting

This section presents common issues that may prevent the switch from connecting to the puppet server.

4.19.7.1 Switch and Server Clocks are not Synchronized

This can be fixed by using NTP to synchronize the clocks at the switch (using the CLI command `ntp`) and at the server (e.g. using `ntpd`).

4.19.7.2 Outdated or Invalid SSL Certificates Either on the Switch or the Server

This can be fixed on the switch using the CLI command `puppet-agent clear-certificates` (requires `puppet-agent restart` to take effect).

On the server it can be fixed by running `puppet cert clean <switch_fqdn>` (FQDN is the Fully Qualified Domain Name which consists of a hostname and a domain suffix).

4.19.7.3 Communications Issue

Make sure it is possible to ping the puppet server hostname from the switch (using the CLI command `ping`).

If the hostname is not reachable (e.g. no DNS server) it can be statically added to the switch local hosts lookup (using the CLI command `ip host`).

Make sure that port 8140 is open (using the command `tracpath {<hostname> | <ip>}/8140`).

4.19.8 Commands

puppet-agent

puppet-agent

Enters puppet agent configuration mode.

Syntax Description	N/A
Default	None
Configuration Mode	Config
History	3.3.4200
Role	admin
Example	switch (config) # puppet-agent switch (config puppet-agent) #
Related Commands	
Notes	

master-hostname

master-hostname <hostname>
no master-hostname

Sets the puppet server hostname.
The no form of the command resets the parameter to its default.

Syntax Description	hostname	Puppet server hostname. Free string may be entered.
Default	puppet	
Configuration Mode	Config Puppet	
History	3.3.4200	
Role	admin	
Example	switch (config puppet-agent) # master-hostname my-puppet-server-host-name switch (config puppet-agent) #	
Related Commands		
Notes		

enable

enable
no enable

Enables the puppet server on the switch.
The no form of the command disables the puppet server.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config Puppet
History	3.3.4200
Role	admin
Example	switch (config puppet-agent) # enable switch (config puppet-agent) #
Related Commands	
Notes	

run-interval

run-interval <time>

Configures the time interval in which the puppet agent reports to the puppet server.

Syntax Description	time	Can be in seconds (“30” or “30s”), minutes (“30m”), hours (“6h”), days (“2d”), or years (“5y”).
Default	30m	
Configuration Mode	Config Puppet	
History	3.3.4302	
Role	admin	
Example	switch (config puppet-agent) # run-interval 40m switch (config puppet-agent) #	
Related Commands	show puppet-agent	
Notes		

restart

puppet-agent restart

Restarts the puppet agent.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config Puppet
History	3.3.4200
Role	admin
Example	switch (config puppet-agent) # restart switch (config puppet-agent) #
Related Commands	
Notes	

show puppet-agent

show puppet-agent

Displays Puppet agent status and configuration.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.4200
	3.3.4302 Updated output with run interval
Role	admin
Example	switch (config puppet-agent) # show puppet-agent Puppet agent is disabled Puppet master hostname: puppet Run interval: 40m switch (config puppet-agent) #
Related Commands	
Notes	

show puppet-agent log

show puppet-agent log **[[not] [matching | continuous] <string> | files [[not] matching] <string>]**

Displays the Puppet agent's log file.

Syntax Description	continuous	Puppet agent log messages as they arrive.
	files	Displays archived Puppet agent log files.
	matching	Displays Puppet agent log that match a given string.
	not	Displays Puppet agent log that do not meet a certain string.
	string	Free string.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.4200	
Role	admin	
Example	<pre>switch (config puppet-agent) # show puppet-agent log Mon Nov 04 11:52:42 +0000 2013 Puppet (notice): Starting Puppet client version 3.2.3 Mon Nov 04 11:52:44 +0000 2013 Puppet (warning): Unable to fetch my node definition, but the agent run will continue: Mon Nov 04 11:52:44 +0000 2013 Puppet (warning): Could not intern from pson: source '"<Puppet::Node:0x7f' not in PSON! Mon Nov 04 11:53:21 +0000 2013 /Netdev_vlan[Vlan104]/ensure (notice): created Mon Nov 04 11:53:22 +0000 2013 /Netdev_vlan[Vlan101]/ensure (notice): created Mon Nov 04 11:53:23 +0000 2013 /Netdev_vlan[Vlan102]/ensure (notice): created Mon Nov 04 11:53:24 +0000 2013 /Netdev_vlan[Vlan103]/ensure (notice): created Mon Nov 04 11:53:40 +0000 2013 /Netdev_l2_interface[ethernet 1/6]/untagged_vlan (notice): untagged_vlan changed 'default' to 'Vlan103' Mon Nov 04 11:53:43 +0000 2013 /Netdev_l2_interface[ethernet 1/7]/untagged_vlan (notice): untagged_vlan changed 'default' to 'Vlan103' Mon Nov 04 11:53:48 +0000 2013 /Netdev_vlan[Vlan100]/ensure (notice): created Mon Nov 04 11:53:48 +0000 2013 /Netdev_l2_interface[ethernet 1/5]/vlan_tagging (notice): vlan_tagging changed 'enable' to 'disable' Mon Nov 04 11:53:48 +0000 2013 /Netdev_l2_interface[ethernet 1/5]/tagged_vlans (notice): tagged_vlans changed '[' to 'Vlan100,Vlan101,Vlan102]' Mon Nov 04 11:53:51 +0000 2013 /Netdev_l2_interface[ethernet 1/1]/tagged_vlans (notice): tagged_vlans changed '[' to '[Vlan101,Vlan104]' Mon Nov 04 11:53:51 +0000 2013 /Netdev_l2_interface[ethernet 1/1]/untagged_vlan (notice): untagged_vlan changed 'default' to 'Vlan100' Mon Nov 04 11:53:54 +0000 2013 /Netdev_l2_interface[ethernet 1/3]/tagged_vlans (notice): tagged_vlans changed '[' to '[Vlan101,Vlan104]' Mon Nov 04 11:53:54 +0000 2013 /Netdev_l2_interface[ethernet 1/3]/untagged_vlan (notice): untagged_vlan changed 'default' to 'Vlan100' Mon Nov 04 11:53:58 +0000 2013 /Netdev_l2_interface[ethernet 1/4]/vlan_tagging (notice): vlan_tagging changed 'enable' to 'disable' Mon Nov 04 11:53:58 +0000 2013 /Netdev_l2_interface[ethernet 1/4]/tagged_vlans (notice): tagged_vlans changed '[' to 'Vlan100,Vlan101,Vlan102]' Mon Nov 04 11:54:03 +0000 2013 /Netdev_l2_interface[ethernet 1/2]/tagged_vlans (notice): tagged_vlans changed '[' to '[Vlan101,Vlan104]' Mon Nov 04 11:54:03 +0000 2013 /Netdev_l2_interface[ethernet 1/2]/untagged_vlan (notice): untagged_vlan changed 'default' to 'Vlan100' Mon Nov 04 11:54:06 +0000 2013 Puppet (notice): Finished catalog run in 47.90 seconds switch (config puppet-agent) #</pre>	
Related Commands		
Notes		

4.20 Virtual Machine

A virtual machine (VM) on a switch is added to allow additional OS to run on top of the switch. The VM OS can connect through mgmt0 interface to the switch system's management interface. In addition, the VM is also connected to the out-of-band network. This allows it to communicate through the network and to control the switch management software.

The number of VMs that may run on a system is user-configurable and also relies on resource availability.



The number of configurable VMs is limited to 4.

Each VM consumes the following resources:

- Memory
- Processing power which is not policed (the user may determine the core to be used)
- MACs which are required for each vNIC (user configurable)

4.20.1 Virtual Machine Configuration

➤ *To configure a VM:*



The example below installs Ubuntu 14 and defines 3GB storage with 512MB memory (default) using the first core of the switch system (default) through mgmt0 interface (default) with an auto-generated MAC (default).

Step 1. Enable the VM feature. Run:

```
switch (config) # virtual-machine enable
```

Step 2. Create a VM. Run:

```
switch (config) # virtual-machine host my-vm
switch (config virtual-machine host my-vm) #
```

Step 3. Define storage for the VM. Run:

```
switch (config virtual-machine host my-vm) # storage create disk size-max 3000
100.0% [#####]
Created empty virtual disk volume 'vdisk001.img' in pool 'default'
Device attached to drive number 1.
switch (config virtual-machine host my-vm) #
```

Step 4. Display the VM parameters (notice boldface). Run:

```
switch (config virtual-machine host my-vm) # show virtual-machine host my-vm
VM 'my-vm'
  Status:      shut off                Architecture:  x86_64
  VCPU used:   0 sec                   Number of VCPUs: 1
  Boot order:  hd, cdrom               Memory size:   512 MB
  Consoles:    text, graphics
  Storage:
    IDE bus, drive 1: default/vdisk001.img (3000 MB capacity)
  Interfaces:
    1: on bridge 'mgmt0'                address unknown  (MAC 52:54:00:2F:89:69)
switch (config virtual-machine host my-vm) # exit
switch (config) #
```

Step 5. Import the VM image. Run:

```
switch (config) # virtual-machine volume fetch url scp://root@<ip>/.../ubuntu-14.04-
server-amd64.iso
Password (if required): *****
100.0% [#####]
```

Step 6. Install the imported image. Run:

```
switch (config) # virtual-machine host my-vm
switch (config virtual-machine host my-vm) # install cdrom file ubuntu-14.04-server-
amd64.iso
```

Step 7. Switch to a different terminal, and run the following command to connect VNC viewer to the VM:

```
$ vncviewer -via admin@<switch IP> 127.0.0.1:0
...
Mellanox MLNX-OS Switch Management

Password: *****
```

Continue VM installation from the VNC prompt.



The switch prompt is unresponsive pending a successful VM installation. Successful VM installation is indicated by the reboot of the VM.



VM IP is determined by DHCP configuration according to the MAC address in [Step 4](#).

➤ **To verify VM configuration, run:**

```
switch (config virtual-machine host my-vm) # show virtual-machine host my-vm
VM 'my-vm'
  Status:      running           Architecture:  x86_64
  VCPU used:   12 min 27.440 sec  Number of VCPUs: 1
  Boot order:  cdrom, hd         Memory size:   512 MB
  Consoles:    text, graphics
  Storage:
    IDE bus, drive 1: default/vdisk001.img (3000 MB capacity)
    IDE bus, drive 2: default/ubuntu-14.04-server-amd64.iso (564 MB capacity) READ-ONLY
  Interfaces:
    1: on bridge 'mgmt0'          address unknown (MAC 52:54:00:2F:89:69)
```

➤ **To perform a VM installation from a USB stick:**



USB stick with supported VM image should be supplied to the user by Mellanox.

- Step 1.** Insert the USB stick (supplied by Mellanox) to the USB port of your switch system.
- Step 2.** Decide on a name for the VM (e.g. “my_vm”).
- Step 3.** Decide on the network configuration of the VM.
 - Use DHCP or alternately use static IP definitions
 - Assign a MAC address or alternately use the default MAC address
- Step 4.** Launch the full installation of the VM with the network definitions of your choice.
For a configuration example, please refer to [Section C.1, “Deploying Mellanox NEO™ on a MLNX-OS® Switch,”](#) on page 1601.

4.20.2 Commands

4.20.2.1 Config

virtual-machine enable

virtual-machine enable
no virtual-machine enable

Enables VM feature on the switch.
The no form of the command disables VM feature on the switch.

Syntax Description	N/A
Default	no virtual-machine enable
Configuration Mode	Config
History	3.4.0000
Role	admin
Example	switch (config) # virtual-machine enable
Related Commands	
Notes	

virtual-machine host

virtual-machine host <vm-name>
no virtual-machine host <vm-name>

Creates a VM, or enters its configuration context if it already exists.
The no form of the command removes the VM with the specified name.

Syntax Description	vm-name	Configures a name for the VM
Default	N/A	
Configuration Mode	Config	
History	3.4.0000	
Role	admin	
Example	switch (config)# virtual-machine host my-vm switch (config virtual-machine host my-vm)#	
Related Commands		
Notes		

arch

arch {i386 | x86_64}

Configures VM CPU architecture.

Syntax Description	i386	32-bit x86 CPU architecture
	x86_64	64-bit x86 CPU architecture
Default	x86_64	
Configuration Mode	Config Virtual Machine Host	
History	3.4.0000	
Role	admin	
Example	switch (config virtual-machine host my-vm)# arch i386	
Related Commands	virtual-machine	
Notes		

comment

comment <string>
no comment

Configures a comment describing the VM.
 The no form of the command deletes the configured comment.

Syntax Description	string	Free string
Default	N/A	
Configuration Mode	Config Virtual Machine Host	
History	3.4.0000	
Role	admin	
Example	switch (config virtual-machine host my-vm)# comment "example VM"	
Related Commands	virtual-machine	
Notes	To configure a multi-word string, the string must be placed within quotation marks.	

console

console {connect [graphics | text [force]] | graphics vnc | text tty}
no console {graphics vnc | text tty}

Configures or connects to a text or graphical console.
 The no form of the command clears console settings.

Syntax Description	connect	Connects to the text console unless specified otherwise: - graphics – connects to the X11 graphical (VNC) console - text – connects to the text console
	graphics vnc	Enables graphical (VNC) console access
	text tty	Enables TTY text console access
Default	Graphical and textual consoles are enabled	
Configuration Mode	Config Virtual Machine Host	
History	3.4.0000	
Role	admin	
Example	switch (config virtual-machine host my-vm)# console connect text	
Related Commands	virtual-machine ssh server x11-forwarding enable	
Notes	- To exit the text console press Ctrl-6 (or Ctrl-Shift-6) - If the guest OS is not configured to receive input from a serial console (ttyS0), the VM console becomes unresponsive when connected to. - To view the graphical console, X display must be enabled. There are two options to activate it, the command <code>vncviewer -via admin@<switchIP> 127.0.0.1:<VNC display num></code> (which is run from an external Linux host) and the command <code>ssh server x11-forwarding enable</code> (which is run from within the switch and requires that you log out and log back in again using <code>ssh -X</code>). The latter command weakens the switch security, therefore, it is recommended to opt for the second option. The VNC display num parameter may be procured by running the command <code>show virtual-machine <vm-name> detail</code>.	

install

install {cancel | cdrom [pool <pool-name>] {file <volume-name> [connect-console <console-type> | disk-overwrite | timeout {<minutes> | none}]}}

Installs an operating system onto this VM (temporarily attach a CD and boot from it).

Syntax Description	cancel	Cancels an install already in progress
	cdrom	Installs an operating system from a CD-ROM (ISO) image
	pool <pool-name>	Configures storage pool in which to find image to install: • default • usb
	file <volume-name>	Specifies CD-ROM (ISO) image from which to install
	connect-console <console-type>	Connects to the console during installation. The types may be: • text – text console • graphics – graphical console
	disk-overwrite	Installs even if primary target volume is not empty
	timeout {<minutes> none}	Configures a timeout for installation in minutes (default is no timeout).
Default	N/A	
Configuration Mode	Config Virtual Machine Host	
History	3.4.0000	
Role	admin	
Example	switch (config virtual-machine host my-vm)# install cdrom pool usb file <image>	
Related Commands	virtual-machine	
Notes	The default pool from which the system installs the ISO image is the /var/ partition in the switch.	

install-from-usb

install-from-usb [ip-address <ip-address> <mask> default-gateway <gw-ip>
[mac <mac-address>] | mac <mac-address>]

Installs a VM including resource allocation and network configurations from a VM image file located on a USB stick.

Syntax Description	ip-address	The IP address to configure for the installed VM
	mask	The IP mask to configure to the installed VM Format example: /24 or 255.255.255.0 Note that a space is required between the IP address and the netmask length
	default-gateway	The IP address of the default gateway to configure for the installed VM
	mac	The MAC address to configure for the installed VM (e.g. ff:ee:dd:cc:bb:aa)
Default	N/A	
Configuration Mode	Config Virtual Machine Host	
History	3.6.2002	
Role	admin	
Example	<pre>switch (config virtual-machine host my-vm)# install-from-usb 100.0% [#####] VM host my-vm MAC is: aa:bb:cc:dd:ee:ff switch (config virtual-machine host my-vm)#</pre>	
Related Commands	virtual-machine	
Notes	USB stick supplied by Mellanox must be inserted into the USB port of the switch system prior to running this command.	

interface

interface <id> {bridge <bridge> | macaddr <mac> | model <model> | name <name>}

Configures virtual interfaces.

Syntax Description	<id>	Interface ID number (1-8 permitted)
	bridge <bridge>	Configures bridge for this interface (i.e. mgmt0 or mgmt1)
	macaddr <mac>	Configures MAC address (e.g. ff:ee:dd:cc:bb:aa)
	model <model>	Configures virtual interface model: • realtek-8139 – Realtek 8139 (default) • virtio – Virtual IO
	name <name>	Configures virtual interface name. The name must begin with “vif”.
Default	N/A	
Configuration Mode	Config Virtual Machine Host	
History	3.4.0000	
Role	admin	
Example	switch (config virtual-machine host my-vm)# interface 1 model virtio	
Related Commands	virtual-machine	
Notes		

memory

memory <MB>

Configures memory allowance.

Syntax Description	MB	Size in megabytes.
Default	512MB	
Configuration Mode	Config Virtual Machine Host	
History	3.4.0000	
Role	admin	
Example	switch (config virtual-machine host my-vm)# memory 1024	
Related Commands	virtual-machine	
Notes	It is recommended not to allocate more than 1GB of memory per VM.	

power

power {cycle [force | connect-console {graphics | text}] | off [force] | on [connect-console {graphics | text}]}

Turns the VM on or off, or other related options.

Syntax Description	cycle	Powers the VM down and then on again immediately
	force	Forces an action on the system.
	connect-console <console-type>	Connects to the console after power-on. The types may be: • text – text console • graphics – graphical console
	off	Powers down the VM
	on	Powers on VM:
Default	N/A	
Configuration Mode	Config Virtual Machine Host	
History	3.4.0000	
Role	admin	
Example	switch (config virtual-machine host my-vm)# power cycle force	
Related Commands	virtual-machine	
Notes		

storage create

storage create disk [drive-number <number> | file <filename> | mode {read-only | read-write} | pool <pool-name> | size-max <MB>]

Creates a new storage device for the VM, with an automatically assigned name.

Syntax Description	create disk	Creates a new virtual disk image for this VM.
	drive-number <number>	Specifies the drive number to be assigned to the volume. Insert “new” to assign a new drive number to the volume.
	file <filename>	Specifies filename for new volume to be created
	mode {read-only read-write}	Specifies initial device mode
	pool <pool-name>	Specifies storage pool in which to create new volume
	size-max <MB>	Specifies maximum disk capacity in megabytes
Default	N/A	
Configuration Mode	Config Virtual Machine Host	
History	3.4.0000	
Role	admin	
Example	switch (config virtual-machine host my-vm)# storage create disk size-max 2000	
Related Commands	virtual-machine	
Notes		

storage device

storage device [bus ide] drive-number <number> [mode {read-only | read-write}] source {[pool <pool-name>] file <filename>}
no storage device [bus ide] drive-number <id>

Modifies existing storage device, or create a new one with a specific name.
The no form of the command removes a storage device from the VM.

Syntax Description	device	Modifies existing storage device, or creates a new one with a specific name
	bus ide	Configures bus type to IDE
	drive-number <number>	Selects device to configure by drive number
	mode {read-only read-write}	Configures the device mode: • read-only – sets the read-only attribute of the volume • read-write – sets the read-write attribute of the volume
	source	Specifies where the data for this volume resides
	file <filename>	Specifies the filename for this volume
	pool <pool-name> file <filename>	Specifies the storage pool for this volume
Default	N/A	
Configuration Mode	Config Virtual Machine Host	
History	3.4.0000	
Role	admin	
Example	switch (config virtual-machine host my-vm)# storage create disk bus ide	
Related Commands	virtual-machine	
Notes		

vcpus

vcpus {count <count> | vcpu <vcpu> pin <cpu-list> [<cpu-list>]}
no vcpus {pin | vcpu <vcpu> pin}

Specifies virtual CPUs.
 The no form of the command removes certain CPU configuration.

Syntax Description	count <count>	Specifies the number of virtual CPUs
	vcpu <vcpu>	Specifies options for a particular virtual CPU
	pin <cpu-list>	Specifies physical CPUs to pin to this vCPU
Default	N/A	
Configuration Mode	Config Virtual Machine Host	
History	3.4.0000	
Role	admin	
Example	switch (config virtual-machine host my-vm)# vcpus count 1	
Related Commands		
Notes		

virtual-machine volume fetch url

```
virt volume fetch url <download-url> [filename <filename> | pool <pool-name>
filename <filename>]
```

Fetches volume image from a remote host.

Syntax Description	download-url	Specifies URL from which to fetch a volume. Format: http, https, ftp, tftp, scp and sftp are supported (e.g. scp://username[:password]@hostname/path/filename)
	filename <filename>	Specifies new filename for fetched volume image
	pool-name <pool-name>	Specifies storage pool for fetched volume image
Default	N/A	
Configuration Mode	Config Virtual Machine Host	
History	3.4.0000	
Role	admin	
Example	switch (config) # virtual-machine volume fetch scp://admin[:admin-pass]@hostname/path/filename>	
Related Commands		
Notes		

virt volume file

virt volume file <name> {create disk size-max <MB> | move {new-name <new-name> | pool <pool-name> new-name <new-name>} | upload <upload-url>}
no virt volume file <volume-name>

Specifies name of volume file to manage.
 The no form of the command deletes the volume file.

Syntax Description	file <name>	Specifies name of volume file to manage
	create	Creates a new volume file under this name
	disk size-max <MB>	Specifies maximum capacity of virtual disk to create
	move	Moves or renames this volume
	new-name <filename>	Specifies a name for the destination file
	pool <pool-name> new-name <filename>	Specifies a storage pool for the copy
	upload <upload-url>	Uploads this volume file to a remote host. Format: ftp, tftp, scp and sftp are supported (e.g. scp://username[:password]@hostname/path/file-name)
Default	N/A	
Configuration Mode	Config Virtual Machine Host	
History	3.4.0000	
Role	admin	
Example	switch (config) # virt volume file my-vm_file create cdrom extract cdrom1	
Related Commands		
Notes		

4.20.2.2 Show

show virtual-machine configured

show virtual-machine configured

Displays global virtualization configuration.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.4.0000
Role	admin
Example	<pre>switch (config) # show virtual-machine configured Virtualization enabled: yes Virtual machines: 2 configured Virtual networks: 0 configured switch (config) #</pre>
Related Commands	
Notes	

show virtual-machine host

show virtual-machine host [<vm-name>]

Displays status for this VM.

Syntax Description	vm-name	The name of the VM.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.4.0000	
Role	admin	
Example	<pre>switch (config) # show virtual-machine host my-vm VM 'my-vm' Status: shut off Architecture: x86_64 VCPU used: 0 sec Number of VCPUs: 1 Boot order: hd, cdrom Memory size: 512 MB Consoles: text, graphics Storage: IDE bus, drive 1: default/vdisk001.img (3000 MB capacity) Interfaces: 1: on bridge 'mgmt0' address unknown (MAC 52:54:00:2F:89:69) switch (config) #</pre>	
Related Commands		
Notes	If the command is run in the middle of an installation, the following banner appears: <pre>*** INSTALL IN PROGRESS: begun <time> ago ***</pre>	

show virtual-machine host configured

show virtual-machine host <vm-name> configured [detail]

Displays configuration for this VM.

Syntax Description	vm-name	The name of the VM.
	detail	Displays detailed configuration for this VM.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.4.0000	
Role	admin	
Example	<pre>switch (config) # show virtual-machine host my-vm configured detail VM 'my-vm' UUID: 0a177a99-f780-5951-877a-bd660e12e5db Text console: enabled Graphics console: enabled Auto-power: last Boot order: hd, cdrom Architecture: x86_64 Memory size: 512 MB Features: ACPI, APIC Number of VCPUs: 1 (No VCPUs pinned) Storage: IDE bus, drive 1 Source pool: default Source file: vdisk001.img (3000 MB capacity) Mode: read-write Interfaces: Interface 1 Name: vif1 MAC address: 52:54:00:2F:89:69 Model: realtek-8139 Bound to: bridge 'mgmt0' switch (config) #</pre>	

Related Commands

Notes

show virtual-machine host detail

show virtual-machine host <vm-name> detail

Displays detailed status for this VM.

Syntax Description	vm-name	The name of the VM.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.4.0000	
Role	admin	
Example	<pre> switch (config) # show virtual-machine host my-vm detail VM 'my-vm' Status: shut off UUID: 0a177a99-f780-5951-877a-bd660e12e5db Text console: enabled Device: N/A Graphics console: enabled VNC display num: N/A Boot order: hd, cdrom Architecture: x86_64 Memory size: 512 MB Features: ACPI, APIC Number of VCPUs: 1 (State of individual VCPUs unavailable when VM is powered off) Storage: IDE bus, drive 1 Source pool: default Source file: vdisk001.img (3000 MB capacity) Mode: read-write Device type: disk Read requests: N/A Read bytes: N/A Write requests: N/A Write bytes: N/A Interfaces: Interface 1 Name: vif1 MAC address: 52:54:00:2F:89:69 Model: realtek-8139 Bound to: bridge 'mgmt0' IP address: RX bytes: 0 RX packets: 0 RX errors: 0 RX drop: 0 TX bytes: 0 TX packets: 0 TX errors: 0 TX drop: 0 switch (config) # </pre>	

Related Commands

Notes

show virtual-machine install

show virtual-machine host <vm-name> install

Displays status of installation of guest OS.

Syntax Description	vm-name	The name of the VM.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.4.0000	
Role	admin	
Example	<pre>switch (config) # show virtual-machine host my-vm install Install status for VM 'my-vm' Install in progress, begun 2 minutes 28 seconds ago. No previous install information available. switch (config) #</pre>	
Related Commands		
Notes		

show virtual-machine interface

show virtual-machine host <vm-name> interface [brief | configure]

Displays full status of all interfaces for this VM.

Syntax Description	vm-name	The name of the VM.
	brief	Displays brief status of all interfaces for this VM.
	configure	Displays configuration of all interfaces for this VM.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.4.0000	
Role	admin	
Example	<pre>switch (config) # show virtual-machine host my-vm interface Interface 1 Name: vif1 MAC address: 52:54:00:2F:89:69 Model: realtek-8139 Bound to: bridge 'mgmt0' IP address: RX bytes: 0 RX packets: 0 RX errors: 0 RX drop: 0 TX bytes: 0 TX packets: 0 TX errors: 0 TX drop: 0 switch (config) #</pre>	
Related Commands		
Notes		

show virtual-machine storage

show virtual-machine host <vm-name> storage

Displays statistics for attached storage.

Syntax Description	vm-name	The name of the VM.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.4.0000	
Role	admin	
Example	<pre>switch (config) # show virtual-machine host my-vm storage Storage for VM 'my-vm' IDE bus, drive 1 Source pool: default Source file: vdisk001.img (3000 MB capacity) Mode: read-write Device type: disk Read requests: N/A Read bytes: N/A Write requests: N/A Write bytes: N/A switch (config) #</pre>	
Related Commands		
Notes		

4.21 Back-Up Battery Units



This section is relevant for the SX6720 switch system.

The back-up battery units (BBUs) are optional field replaceable Lithium-Ion batteries that provide backup power in the event the main AC power supply is disrupted. The batteries are guaranteed to provide up to 5 minutes of stand-by power that allows the user to continue operations or alternately to perform any kind of backup or recovery operations. However, depending on the system load, the batteries may provide power for more time.

The BBU provides power outage for a limited time which allows time for performing controlled shutdown or recovering the main power source of the system, or connecting to a new power source.

When the switch system is plugged in to a power source, the BBU may be configured to charge in preparation for possible power outage scenarios.



Switching from the main power source to the BBU and vice versa does not interfere with the normal operation of the switch system.



BBU discharge must be disabled before any planned shutdown of the switch.

4.21.1 BBU Calibration Procedure

During normal operation, the BBU's gas gauge parameters are not updated and its accuracy is degraded. The BBU can be calibrated, thereby adjusting the Q_{max} and RA tables. BBU calibration can be performed using the command "battery-backup-unit calibrate-battery [foreground]".

During calibration, the battery under calibration is completely discharged, analyzed and recharged. The calibration process requires that both BBUs are fully charged before calibration begins. This ensures that the battery not being calibrated can support the platform in case of power failure during the calibration process.

To maintain the battery calibration, it is recommended that calibration is performed every 3 months or when the Max Error value (shown in the command "show battery-backup-unit details") becomes equal to or greater than 3%.

The last calibration result is saved in the switch and can be viewed using the command "show battery-backup-unit details". If the switch is reset or the BBU is ejected, this information is reset and the last calibration date and calibration results become "N/A".

The value of "Last Calibration" in the output of the command "show battery-backup-unit details" may be one of the following:

- PASSED & date – last calibration has completed successfully

- FAILED – last calibration has failed
- In-progress – BBU is currently being calibrated
- N/A – no valid information on this BBU’s last calibration result. This is displayed when the BBU has not been since last insertion.



The BBU must be fully charged for at least 3 hours before calibrating it.

4.21.2 BBU Self-Test

The BBU has a self-testing procedure which allows the user to verify that the BBU is operative and capable to provide power to the platform if required. The procedure lasts for 10 seconds.

In this procedure the power supply’s (PS) voltage is lowered causing the BBU to become the system’s primary power source. If the BBU is faulty, the platform will still get power from the PS.

The test is triggered manually by the user with the command “battery-backup-unit test-battery” and the user is then prompted on whether the BBU has passed or failed the test.



The test is run only if the BBU is fully charged.

During the test, the BBU remains operative so in case of power failure the BBU is able to keep the system running, and the test is completed regularly.

4.21.3 BBU Shut-Off Timer

The BBU shut-off timer allows the user to configure a finite time after which the system definitely shuts off when on battery power. The timer starts when the AC power supply is lost to all batteries and the entire system is running purely on BBU power. When the timer expires, the system will disable the BBUs and immediately shut down. Whenever AC power is restored, the system reboots and resumes normal operation.



A prerequisite for the timer functionality to be effective is that the batteries are “discharge enabled”. Otherwise, the system shuts down immediately after AC power is lost.

The timer can be configured with a timeout duration ranging from 1 to 60 minutes. After an upgrade from a software image that does not have this feature, the timer is disabled by default. The user must explicitly configure a timer duration in order to enable it. Configuring the timer or changing its settings is only permitted when the system is on AC power.

When the timer is enabled, the system displays an estimated battery runtime which represents the amount of time the system will continue to run on battery power.

With the shut-off timer enabled, the duration displayed will be lowest of the following values:

- 5 minutes

- Remaining timer duration
- Calculated remaining battery capacity

Assume, for example, that based on the current load, the battery can provide power for around 30 minutes. If the shut-off timer is configured for 10 minutes, the system will show 5 minutes because that is the lowest among the three aforementioned values. As the system continues to operate on battery power, and the timer continues to count down, at some point the remaining timer duration will become less than 5 minutes. At that point the system shall show the remaining timer duration (assuming the battery capacity is larger). This information is also displayed on the WebUI, however the timer can only be set via the CLI.

4.21.4 Commands

battery-backup-unit charge

battery-backup-unit charge {bbu1 | bbu2 | all}

no battery-backup-unit charge {bbu1 | bbu2 | all}

Enables charging the BBU.

The no form of the command disables charging the BBU.

Syntax Description	bbu1	Enables charging of BBU1
	bbu2	Enables charging of BBU2
	all	Enables charging of both BBUs
Default	Enabled	
Configuration Mode	Config	
History	3.4.1100	
Role	admin	
Example	switch (config) # battery-backup-unit charge all	
Related Commands	N/A	
Notes	- When disabled, the BBU protection against power outage is still available, however the BBU will not be charged - For example, this command can be used when there is a need to lower the chassis power consumption	

battery-backup-unit discharge

battery-backup-unit discharge {bbu1 | bbu2 | all}
no battery-backup-unit discharge {bbu1 | bbu2 | all}

Enables discharging the BBU in case of power outage.
The no form of the command disables discharging the BBU in case of power outage.

Syntax Description	bbu1	Enables charging of BBU1
	bbu2	Enables charging of BBU2
	all	Enables charging of both BBUs
Default	Enabled	
Configuration Mode	Config	
History	3.4.1100	
Role	admin	
Example	switch (config) # battery-backup-unit discharge all	
Related Commands	N/A	
Notes	- Warning: When disabled, the BBU does not provide protection against power outage - This command must be run before unplugging the switch system from the main. Otherwise, it shall continue to run on BBU power. - BBU discharge must be disabled before any planned shutdown of the switch.	

battery-backup-unit calibrate-battery

battery-backup-unit calibrate-battery {bbu1 | bbu2} [force]
no battery-backup-unit calibrate-battery {bbu1 | bbu2}

Starts BBU calibration process in the background.
 The no form of the command cancels the active calibration process.

Syntax Description	bbu1	Calibrates BBU1
	bbu2	Calibrates BBU2
	force	Starts the calibration process even if the other BBU is not fully charged
Default	N/A	
Configuration Mode	Config	
History	3.4.1854	
Role	admin	
Example	switch (config) # battery-backup-unit calibrate-battery bbu1	
Related Commands	show battery-backup-unit show battery-backup-unit details	
Notes	• BBU must be fully charged for at least 3 hours before calibrating it • When checking the status of the BBUs, if under calibration, the charging state indicates under which stage of calibration the BBU is at the moment • BBU calibration may take several hours to complete • If BBU calibration is canceled, the status of the last calibration in the command “show back-up-battery details” appears as failed	

battery-backup-unit calibrate-battery foreground

battery-backup-unit calibrate-battery {bbu1 | bbu2} foreground [force]

Starts BBU calibration process in the foreground.

Syntax Description	bbu1	Calibrates BBU1
	bbu2	Calibrates BBU2
	force	Starts the calibration process even if the backup BBU is not fully charged
Default	N/A	
Configuration Mode	Config	
History	3.4.1854	
Role	admin	
Example	<pre>switch (config) # battery-backup-unit calibrate-battery bbu2 foreground BBU2 Calibration: remaining capacity 0% ##### BBU2 Calibration: finished discharge phase. Waiting for the battery to relax BBU2 Calibration: relaxation phase (56 min passed) ##### Calibration Result: PASSED switch (config) #</pre>	
Related Commands	N/A	
Notes	• BBU must be fully charged for at least 3 hours before calibrating it • When checking the status of the BBUs, if under calibration, the charging state indicates under which stage of calibration the BBU is at the moment • Battery calibration may take several hours to complete • To send the calibration process to the background, use the key combination Ctrl+C • To cancel the calibration process, send the calibration to the background and then run the command “no battery-backup-unit calibrate-battery {bbu1 bbu2}”	

battery-backup-unit shut-off-timer

battery-backup-unit shut-off-timer <time>
no battery-backup-unit shut-off-timer

Enables and configures BBU shut-off timer.
The no form of the command disables BBU shut-off timer.

Syntax Description	time	Timer in minutes Range: 1-60
Default	Disabled	
Configuration Mode	Config	
History	3.6.1002	
Role	admin	
Example	switch (config) # battery-backup-unit shut-off-timer 10	
Related Commands	N/A	
Notes		

battery-backup-unit test-battery

battery-backup-unit test-battery {bbu1 | bbu2}

Starts BBU self-test.

Syntax Description	bbu1	Tests BBU1
	bbu2	Tests BBU2
Default	N/A	
Configuration Mode	Config	
History	3.4.2008	
Role	admin	
Example	<pre>switch (config) # battery-backup-unit test-battery bbu1 BBU[1] self-test PASSED switch (config) #</pre>	
Related Commands	N/A	
Notes	• CLI hangs for 15 seconds while the test is running • BBU under test must be fully charged before running test	

show battery-backup-unit

show battery-backup-unit

Displays the present BBU devices.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.4.1100 First version 3.4.1854 Updated Example output 3.6.1002 Updated Example output with BBU Shut-off Timer
Role	admin
Example	<pre>switch (config) # show battery-backup-unit ----- - Module Device Charging state Charge Discharge ----- - bbu1 controller Discharging 100%, remaining 3:49 M enable enable bbu2 controller Discharging 100%, remaining 3:49 M enable enable BBU Shut-off Timer : 10:00 minutes Estimated battery run-time : 03:49 minutes</pre>
Related Commands	N/A
Notes	The possible charging states are as follows: • Under Calibration: discharging – the BBU is in the discharging stage of the calibration • Under Calibration: relaxing – the BBU is in the relaxing stage of the calibration • Under self-test – the BBU is undergoing self-test • Charging – the BBU is being charged • Discharging – the BBU is being discharged • Fully charged – the BBU is fully charged • Fully discharged – the BBU is depleted • Not charging – the BBU is neither charging nor discharging due to charge/discharge configuration

show battery-backup-unit details

show battery-backup-unit {bbu1 | bbu2 | all} details

Displays the present BBU devices.

Syntax Description	bbu1	Displays details of BBU1
	bbu2	Displays details of BBU2
	all	Displays details of both BBUs
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.4.1100	First version
	3.4.1138-01	Added firmware version
	3.4.1854	Updated Example
	3.5.1000	Updated Example
Role	admin	
Example	<pre>switch (config) # show battery-backup-unit bbu1 details battery backup unit [bbu1] details Charging state : Fully charged Charge rate : fast Alarms : none Relative capacity : 83360 mWh / 100% Absolute capacity : 83580 mWh Designed capacity : 88560 mWh Absolute charge : 95% Manufacture date : 14-39 Serial number : MT1443B01059 Chemistry : Li-Ion Temperature : 24.40 C Voltage : 12.500 V Current : 0 mA BBU FW version : 703 Max error : 1% Last calibration : N/A Last test : N/A PS FW version : 503 switch (config) #</pre>	
Related Commands	N/A	
Notes		

4.22 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.

4.22.1 Configuring IP Table Filtering

Prerequisite for IPv6:

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

➤ *To configure IPv4 table filtering:*

Step 1. Select the policy that applies to the input/output chain. (Default policy is accept.) Run:

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

Step 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
```

Step 3. Enable IP table filtering. Run:

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

Step 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 : -
```

4.22.2 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
```

4.22.3 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.

4.22.4 Commands

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	Disabled
Configuration Mode	Config
History	3.5.1000
Role	admin
Example	switch (config) # ip filter enable switch (config) #
Related Commands	N/A
Notes	It is recommended to run this command only after configuring all of the IP table filter parameters.

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	
Role	admin	
Example	switch (config) # ipv6 filter chain input policy accept switch (config) #	
Related Commands	N/A	
Notes		

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
Role	admin
Example	<pre>switch (config) # ipv6 filter enable chain input rule append tail target drop state related protocol all dup-delete switch (config) #</pre>

Related Commands

N/A

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.
-

4.23 Resource Scale

MLNX-OS 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, IP Proxy-ARP) may use internal resources according to configured allocation policy mode which could be one of the following:

- Loose – a configuration that supports flexible user experience while providing protection to assure some protection against flooding of ARP
- Strict – allows backward compatibility



Loose mode is supported for Ethernet and VPI system profiles only.



Loose mode is supported on PPC systems only in Proxy-ARP mode. Therefore, “no ip proxy-arp” can only be configured in strict mode.



Transition between modes saves configuration and reloads the system.

The default configuration on different types of systems is as follows:

- SwitchX®-2 PPC: Strict
- SwitchX®-2 x86: Strict in VPI profile; Loose in Ethernet profile
- Spectrum™: Loose

4.23.1 Ethernet Resources

4.23.1.1 Strict Mode

Table 50 presents the number of resources available for each SwitchX® based node in strict mode.

Table 50 - Number of Resources per Node in Strict Mode

Resource	Max Resources
Number of ACL rules	1488
Number of MAC addresses	48K
Number of IPv4 neighbors	2048
Number of IPv4 UC routes	4094

Table 50 - Number of Resources per Node in Strict Mode

Resource	Max Resources
Number of IPv4 MC routes	671
Number of IPv4 (ECMP) UC routes	2K

4.23.1.2 Loose Mode – SwitchX®

Table 51 presents the number of resources available for each SwitchX® based node in loose mode.

Table 51 - Number of Resources per Node in Strict Mode for SwitchX Based Systems

Resource	Max Resources
Number of ACL rules	5120
Number of MAC addresses	48K
Number of IPv4 neighbors	8000
Number of IPv4 UC routes	10936
Number of IPv4 MC routes	2047
Number of IPv4 (ECMP) UC routes	5312

4.23.1.3 Loose Mode – Spectrum™

Table 52 presents the number of resources available for each Spectrum™ based node in strict mode.

Table 52 - Number of Resources per Node in Strict Mode for Spectrum Based Systems

Resource	Max Resources
Number of ACL rules	5120
Number of MAC addresses	88K
Number of IPv4 neighbors	8000
Number of IPv4 UC routes	10936
Number of IPv4 MC routes	2047
Number of IPv4 (ECMP) UC routes	5312

4.23.2 Proxy-ARP Resources

4.23.2.1 Strict Mode – Unicast

Table 53 presents the number of resources available for each node in strict Proxy-ARP unicast mode.

Table 53 - Number of Resources per Node in Strict Proxy-ARP Unicast Mode

Resource	Max Resources
Number of ACL rules	1488
Number of UC routes	160
Number of MC routes	0
Number of Ethernet ARP entries	512
Number of InfiniBand ARP entries	3520
Number of InfiniBand GRH ARP entries	0

4.23.2.2 Strict Mode – Multicast

Table 54 presents the number of resources available for each node in strict Proxy-ARP multicast mode.

Table 54 - Number of Resources per Node in Strict Proxy-ARP Multicast Mode

Resource	Max Resources
Number of ACL rules	1488
Number of UC routes	160
Number of MC routes	2500
Number of Ethernet ARP entries	128
Number of InfiniBand ARP entries	1024
Number of InfiniBand GRH ARP entries	0

4.23.2.3 Loose Mode – Unicast

Table 55 presents the maximum number of resources available for each node in Proxy-ARP loose unicast mode.

Table 55 - Maximum Number of Resources per Node in Proxy-ARP Loose Unicast Mode

Resource	Max Resources
Number of ACL rules	5120 in x86
	1488 in PPC
Number of UC routes	160
Number of MC routes	N/A

Table 55 - Maximum Number of Resources per Node in Proxy-ARP Loose Unicast Mode

Resource	Max Resources
Number of Ethernet ARP entries	7500
Number of InfiniBand ARP entries	3750
Number of InfiniBand GRH ARP entries	1875

4.23.2.4 Loose Mode – Multicast

Table 56 presents the maximum number of resources available for each node in Proxy-ARP loose multicast mode.

Table 56 - Maximum Number of Resources per Node in Proxy-ARP Loose Multicast Mode

Resource	Max Resources
Number of ACL rules	5120 in x86
	1488 in PPC
Number of UC routes	160
Number of MC routes	3750
Number of Ethernet ARP entries	7500
Number of InfiniBand ARP entries	3750
Number of InfiniBand GRH ARP entries	1875

4.23.3 Commands

system resource table

system resource table {loose | strict}
no system resource table

Configures system resource table.
The no form of the command restores the system to its default mode.

Syntax Description	loose	Sets system resource table mode as loose
	strict	Sets system resource table mode as strict
Default	SwitchX®-2 PPC: Strict SwitchX®-2 x86: Strict in VPI profile; Loose in Ethernet profile Spectrum™: Loose	
Configuration Mode	Config	
History	3.5.1000	
	3.6.3004	Added support for Proxy-ARP
Role	admin	
Example	switch (config) # system resource table strict	
Related Commands	N/A	
Notes	• x86 based systems support strict and loose modes • Spectrum based systems only support loose mode • Loose mode is supported on PPC systems only in Proxy-ARP mode • Transition between modes saves configuration and reloads the system (after user approval)	

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
	3.6.3004 Added support for Proxy-ARP
Role	admin

Example

```

/**example without Proxy-ARP//

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

switch (config) # show system resource table acl
-----
Table-Id                      In-Use
-----
acl                            0      0
eth-ipv4-uc                    7500   0
ib-ipv4-uc                     3750   0
ib-ipv4-grh                    1875   0
uc-route                       160     0
ipv4-mc                        0       0

Mode: loose
Total configured entries: 0

/**example for Proxy-ARP//

switch (config) # show system resource table
-----
Table-Id                      Total    In-Use
-----
acl                            1504    0
eth-ipv4-uc                    512     0
ib-ipv4-uc                     3520    0
ib-ipv4-grh                     0       0
uc-route                       160     0
ipv4-mc                        0       0

Mode: strict
Total configured entries: 0
Total free entries: 5696

```

Related Commands N/A

Notes

5 Ethernet Switching

5.1 Interface

Interface Ethernet 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 40GbE QSFP interfaces as 10GbE (via QSA adapter), the speed must be manually set with the command “speed 10000” under the interface configuration mode.

5.1.1 Break-Out Cables

The break-out cable is a unique Mellanox capability, where a single physical quad-lane QSFP port is divided into 2 dual-lane ports or 4 single-lane ports. It maximizes the flexibility of the end user to use the Mellanox switch with a combination of dual-lane, single-lane and quad-lane interfaces according to the specific requirements of its network. Certain ports cannot be split at all, and there are ports which can be split into 2 ports only (for more information please refer to your Switch System 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 single-lane 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.

Figure 20: Break-Out Cable



A split-4 operation results in blocking a quad-lane port in addition 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 (four lanes) one of the other ports on the switch is disabled (unmapped).

To see the exact splitting options available per system, refer to each specific system's hardware user manual (Cabling chapter) located on the Mellanox website.

5.1.1.1 Changing the Module Type to a Split Mode

➤ *To split an interface:*

Step 1. Shut down all the ports related to the interface. Run:

- in case of split-2, shut down the current interface only
- in case of split-4, shut down the current interface and the other interface according switch system's spec

```
switch (config) # interface ethernet 1/1
switch (config interface ethernet 1/1) # shutdown
switch (config interface ethernet 1/1) # exit
switch (config) # interface ethernet 1/4
switch (config interface ethernet 1/4) # shutdown
```

Step 2. Split the ports as desired. Run:

```
switch (config interface ethernet 1/4) # module-type qsfp-split-4
switch (config interface ethernet 1/4) #
```

Step 3. The following warning will be displayed:

the following interfaces will be unmapped: 1/420 1/19.

Choose Yes when prompted Type 'yes' to confirm split

The <ports> field in the warning refers to the affected ports from splitting port <inf> in the applied command.



Please beware that splitting a port into 4 prevents you from accessing the splittable port, and an additional one. For example, in the procedure above, ports 3 and 4 become inaccessible.

5.1.1.2 Unsplitting a Split Port

➤ **To unsplit a split port:**

Step 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
```

Step 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) # module-type qsfp
```



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.

Step 3. Type “yes” when prompted “Type 'yes' to confirm unsplit.”

5.1.2 56GbE Link Speed

Mellanox offers proprietary speed of 56Gb/s per Ethernet interface.



The following OPNs support 56GbE:

- MSX6036F-xxxx
- MSX1036x-xxxS
- MSX1024x-xxxS
- MSX1012x-xxxx
- MSX6012F-xxxx
- MSX6018F-xxxx

The following OPNs do not support 56GbE:

- MSX6036T-xxxx
- MSX1036x-xxxR
- MSX6012T-xxxx
- MSX6018T-xxxx



56GbE speed is not supported on SwitchX® (A1) ASIC based switch systems.

➤ **To achieve 56GbE link speed:**

Step 1. Make sure your system is 56Gb/s capable (e.g. SX6036F, SX1024, and SX1036).

Step 2. Install GW license if necessary. Run:

```
switch (config) # license install <license key>
```



For a list of the available licenses see [Section 2.4, “Licenses,”](#) on page 63.

Step 3. Set the system profile to be eth-single-switch, and reset the system:

```
switch (config) # system profile eth-single-profile
```

Step 4. Set the speed for the desired interface to 56GbE as follows. Run:

```
switch (config) # interface ethernet 1/1
switch (config interface ethernet 1/1) # speed 56000
switch (config interface ethernet 1/1) #
```

Step 5. Verify the speed is 56GbE

```
switch (config) # show interfaces ethernet 1/1
Eth1/1
Admin state: Enabled
Operational state: Down
Description: N/A
Mac address: 00:02:c9:5d:e0:26
MTU: 1522 bytes
Flow-control: receive off send off
Actual speed: 56 Gbps
Switchport mode: access
Rx
0 frames
0 unicast frames
0 multicast frames
0 broadcast frames
0 octets
0 error frames
0 discard frames
Tx
0 frames
0 unicast frames
0 multicast frames
0 broadcast frames
0 octets
0 discard frames
switch (config) #
```

5.1.3 Transceiver Information

MLNX-OS 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 1/60 transceiver
Port 1/60 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
switch (config) #
```



The indicated cable length is rounded up to the nearest natural number.

5.1.4 High Power Transceivers

Mellanox switch systems offer high power transceiver (LR4) support in the following ports:

- SX1036/SX1700 – ports 1, 3, 33, 35
- SX1024/SX1400 – ports 50, 52, 54, 56, 58, 60
- SX1012/SX1710 – all ports

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.

5.1.5 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 of errors that can be corrected depends on the FEC algorithm and the amount of redundant data.

100GbE Mellanox-to-Mellanox Ethernet connections always enable standard Reed Solomon (RS) FEC on all cables.

If a Mellanox system is connected to a 3rd party system, then FEC is only activated if the 3rd party requests it also.

5.1.6 Commands

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 in case of split port.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	First version
	3.2.1100	Added range support
Role	admin	
Example	<pre>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) #</pre>	
Related Commands	show interfaces ethernet	
Note		

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	
Role	admin	
Example	switch (config interface ethernet 1/1) # boot-delay 60	
Related Commands	show interfaces ethernet	
Note	This command delays the interface from boot time of the interface Configuration save and system reboot is required for the configuration to take effect.	

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 command implementation.
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 MLAG port-channel configuration mode
Role	admin	
Example	switch (config interface ethernet 1/1) # flowcontrol receive off switch (config interface ethernet 1/1) #	
Related Commands	show interfaces ethernet	
Note	N/A	

module-type

module-type <type> [force]
no module-type <type> [force]

Splits the interface to two or four separate interfaces, or merges them back to a single interface (QSFP).

The no form of the command resets the interface to its default configuration.

Syntax Description	type	qsfp - Port runs at 40000/56000Mbps qsfp-split-2 - Port is split and runs at 2X10000Mb/s qsfp-split-4 - Port is split and runs at 4X10000Mb/s
	force	Force the split operation without asking for user confirmation.
Default	QSFP	
Configuration Mode	Config Interface Ethernet	
History	3.1.1400	
	3.5.0000	
Role	admin	
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 switch (config interface ethernet 1/4) #	
Related Commands		
Note	• The affected interfaces should be disabled prior to the operation • In order to unsplit the interface, use the command with “qsfp”, the speed is set to 40Gb/s “module-type qsfp” • The following speeds are supported on the different Ethernet interface types: • qsfp - 1G, 10G, 25G, 40G, 50G, 56G, 100G • qsfp-split-2 - 1G, 10G, 25G, 50G • qsfp-split-4 - 1G, 10G, 25G	

mtu

mtu <frame-size>

Configures the Maximum Transmission Unit (MTU) frame size for the interface.

Syntax Description	frame-size	This value may be 1500-9216 bytes
Default	1500 bytes	
Configuration Mode	Config Interface Ethernet Config Interface Port Channel Config Interface MLAG Port Channel	
History	3.1.0000	
	3.3.4500	Added MLAG port-channel configuration mode
Role	admin	
Example	switch (config interface ethernet 1/1) # mtu 9216 switch (config interface ethernet 1/1) #	
Related Commands	show interfaces ethernet	
Note		

shutdown

shutdown
no shutdown

Disables the interface.
The no form of the command enables the interface.

Syntax Description	N/A
Default	The 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 MLAG port-channel configuration mode
Role	admin
Example	switch (config interface ethernet 1/1) # shutdown switch (config interface ethernet 1/1) #
Related Commands	show interfaces ethernet
Note	

description

description <string>
no description

Sets 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 MLAG port-channel configuration mode
Role	admin	
Example	switch (config interface ethernet 1/1) # description my-interface switch (config interface ethernet 1/1) #	
Related Commands	show interfaces ethernet	
Note		

speed

speed <port speed> [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	port speed	The following options are available: 1G or 1000 - 1GbE 10G or 10000 - 10GbE 25G or 25000 - 25GbE 40G or 40000 - 40GbE 50G or 50000 - 50GbE 56G or 56000 - 56GbE 100G or 100000 - 100GbE
	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	
	3.3.4500	Added MLAG port-channel configuration mode
	3.5.0000	Added 25GbE, 50GbE, and 100GbE speeds and updated notes
Role	admin	
Example	switch (config interface ethernet 1/1) # speed 40G switch (config interface ethernet 1/1) #	
Related Commands	show interfaces ethernet	
Note	- The default speed depends on the interface capabilities, interface capable of 40GbE will have 40GbE speed by default - SwitchX systems do not support 25GbE, 50GbE, and 100GbE speeds - Not all interfaces support all speed options	

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 MLAG port-channel configuration mode
Role	admin	
Example	switch (config interface ethernet 1/1) # load-interval 30 switch (config interface ethernet 1/1) #	
Related Commands	show interfaces ethernet	
Note	This interval is used for the ingress rate and egress rate counters.	

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
Role	admin
Example	switch (config interface ethernet 1/1) # ip address dhcp switch (config interface ethernet 1/1) #
Related Commands	interface ethernet show interfaces ethernet
Note	

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	• auto – auto-FEC selection • no-fec – disables FEC • fec-on – enables FEC
	force	Forces configuration (does not require toggling interface to take effect)
Default	Auto-FEC selection	
Configuration Mode	Config Interface Ethernet	
History	3.5.0000	
	3.6.2002	Added force option
Role	admin	
Example	switch (config interface ethernet 1/2) # fec-override fec-on	
Related Commands	show interfaces ethernet	
Notes	• This command is supported only on Spectrum™ based switch systems • Use this command with caution. There is no limitation in configuring non-standard FEC. It may cause the link to malfunction.	

clear counters

clear counters

Clears the interface counters.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config Interface Ethernet Config Interface Port Channel
History	3.1.0000 3.3.4500 Added MLAG port-channel configuration mode
Role	admin
Example	switch (config interface ethernet 1/1) # clear counters
Related Commands	show interfaces ethernet
Note	

show interfaces ethernet

show interfaces ethernet <inf>

Displays the configuration and status for the interface.

Syntax Description	inf	Interface number: <slot>/<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
Role	admin	

Example

```
switch (config) # show interfaces ethernet 1/14

Eth1/14
  Admin state: Enabled
  Operational state: Up
  Last change in operational status: 4w4d and 22:35:26 ago (1 oper
change)
  Boot delay time: 60 sec
  Boot delay timer status: N/A
  Description: N/A
  Mac address: f4:52:14:5c:73:f8
  MTU: 1500 bytes(Maximum packet size 1522 bytes)
  Fec: auto
  Flow-control: receive off send off
  Actual speed: 40 Gbps
  Width reduction mode: disabled
  DHCP client: Disabled
  IP Address: 8.9.14.9 /24
  Broadcast address: 8.9.14.255
  Arp timeout: 1500 seconds
  VRF: default
  MAC learning mode: Enabled
  Isolation group: N/A
  Last clearing of "show interface" counters : Never
  60 seconds ingress rate: 168 bits/sec, 21 bytes/sec, 1 packets/sec
  60 seconds egress rate: 160 bits/sec, 20 bytes/sec, 1 packets/sec

Rx
  559480          packets
  4335            unicast packets
  550812          multicast packets
  4333            broadcast packets
  56941600        bytes
  0              error packets
  0              discard packets

Tx
  557579          packets
  4332            unicast packets
  548912          multicast packets
  4335            broadcast packets
  54615032        bytes
  0              error packets
  0              discard packets
```

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 [Section 5.1.4, “High Power Transceivers,”](#) on page 665.

show interfaces ethernet capabilities

show interfaces ethernet [<inf>] capabilities

Displays the interface capabilities.

Syntax Description	inf	Interface number: <slot>/<port>.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # show interfaces ethernet 1/1 capabilities Eth1/1 Speed : 10000,40000 FlowControl : Send, Receive switch (config) #</pre>	
Related Commands		
Note		

show interfaces ethernet counters

show interfaces ethernet <inf> counters [priority <prio>]

Displays the extended counters for the interface.

Syntax Description	inf	Interface number: <slot>/<port>
	priority	Displays interface extended counters per priority. Range: 0-7 or “all”
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.1.0000	
	3.5.1000	Added notes
	3.6.1002	Added “error packets” counter to Tx
Role	admin	

Example

```
switch (config) # show interfaces ethernet 1/1 counters
```

```
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
0      packets
0      unicast packets
0      multicast packets
0      broadcast packets
0      bytes
0      error packets
0      discard packets
0      pause packets
0      TX wait
0      TX wait useconds
0      queue depth TC0
0      queue depth TC1
0      queue depth TC2
0      queue depth TC3
0      queue depth TC4
0      queue depth TC5
0      queue depth TC6
0      queue depth TC7
```

```
switch (config) #
```

Related Commands

Note

- Spectrum™ based systems display queue depth for TC0 - TC7
- SwitchX® based systems display queue depth for TC0 - TC3 only

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
Role	admin	
Example	<pre>switch (config) # show interfaces ethernet description Interface Admin state Operational state ----- Eth1/58 Enabled Down Eth1/59 Enabled Up Eth1/60 Enabled Down (Suspend) switch (config) # show interfaces ethernet 1/60 description Eth1/60 Admin state: Enabled Operational state: Down (Suspend) switch (config) #</pre>	
Related Commands		
Note		

show interfaces ethernet rates

show interfaces ethernet [<inf>] rates [<transfer-rate-unit>]

Displays the current transfer rate of the interface.

Syntax Description	transfer-rate-unit	• 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 readable unit (KB/MB/GB/BS) depending on the range																														
Default	N/A																															
Configuration Mode	Any Command Mode																															
History	3.6.2002																															
Role	admin																															
Example	<pre>switch (config) # show interfaces ethernet rates KB</pre> <table><thead><tr><th>Port</th><th colspan="2">egress</th><th colspan="2">ingress</th></tr><tr><th></th><th>avg rate (KB/s)</th><th>pkts/sec</th><th>avg rate (KB/s)</th><th>pkts/sec</th></tr></thead><tbody><tr><td>Eth1/1</td><td>0</td><td>0</td><td>0.032</td><td>1</td></tr><tr><td>Eth1/2</td><td>0</td><td>0</td><td>0.032</td><td>1</td></tr><tr><td>Eth1/3</td><td>0</td><td>0</td><td>0</td><td>0</td></tr><tr><td>...</td><td></td><td></td><td></td><td></td></tr></tbody></table> <pre>switch (config) #</pre>		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	...				
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																																
Note																																

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		
Role	admin			
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
switch (config) #				
Related Commands				
Note				



show interfaces ethernet transceiver

show interfaces ethernet [*<inf>*] transceiver

Displays the transceiver info.

Syntax Description	inf interface number: <slot>/<port>
Default	N/A
Configuration Mode	Any Command Mode
History	3.1.0000
Role	admin
Example	<pre>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 switch (config) #</pre>
Related Commands	
Note	- For a full list of the supported cables and transceivers, please refer to the LinkX™ Cables and Transceivers webpage in Mellanox.com: http://www.mellanox.com/page/cables?mtag=cable_overview. - If a high power transceiver (e.g. LR4) is used, it will be indicated in the field “cable/module type”.

show interfaces ethernet transceiver counters

show interfaces ethernet [<inf>] transceiver counters

Displays PHY counters.

Syntax Description	inf	interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.6.1002	
Role	admin	
Example	<pre>switch (config) # show interfaces ethernet 1/1 transceiver counters Rx phy received bits 17725862707200 phy symbol errors 0 phy corrected bits 0</pre>	
Related Commands		
Note	- 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) - This command is only supported on Spectrum™ based switch systems	

show interfaces ethernet transceiver counters details

show interfaces ethernet [<inf>] transceiver counters

Displays all PHY counters.

Syntax Description	inf	interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.6.1002	
Role	admin	
Example	<pre>switch (config) # show interfaces ethernet 1/1 transceiver counters details ----- Phy counters ----- Symbol errors 0 Sync headers errors 0 Edpl/bip errors lane0 0 Edpl/bip errors lane1 0 Edpl/bip errors lane2 0 Edpl/bip errors lane3 0 FC corrected blocks lane0 0 FC corrected blocks lane1 0 FC corrected blocks lane2 0 FC corrected blocks lane3 0 FC uncorrectable blocks lane0 0 FC uncorrectable blocks lane1 0 FC uncorrectable blocks lane2 0 FC uncorrectable blocks lane3 0 RS corrected blocks 0 RS uncorrectable blocks 0 RS no errors blocks 1130552748 RS single error blocks 0 RS corrected symbols total 0 RS corrected symbols lane0 0 RS corrected symbols lane1 0 RS corrected symbols lane2 0 RS corrected symbols lane3 0 Link down events 0 Successful recovery events 0 Time since last clear 176127</pre>	
Related Commands		
Note	The number of lanes displayed depends on interface splitter ratio (4-way-split – each split has only 1 lane; 2-way-split – each split has 2 lanes)	

show interfaces ethernet transceiver diagnostics

show interfaces ethernet [*<inf>*] transceiver diagnostics

Displays cable channel monitoring and diagnostics info for this interface.

Syntax Description	inf	Interface number: <i><slot>/<port></i>
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.6.2002	
Role	admin	

Example

```
switch (config) # show interfaces ethernet 1/1 transceiver diagnostics
Port 1/1 transceiver diagnostic data:
```

```

Temperature (-127C to +127C)
  Temperature           : 33 C
  Hi Temp Alarm Thresh  : 17 C
  Low Temp Alarm Thresh : 2 C
  Temperature Alarm     : None

Voltage ( 0 to 6.5535 V)
  Voltage               : 3.29450 V
  Hi Volt Alarm Thresh  : 3.70000 V
  Low Volt Alarm Thresh : 2.90000 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 to 6.5535 mW)
  Ch1 Tx Power          : 1.03080 mW
  Ch2 Tx Power          : 1.05070 mW
  Ch3 Tx Power          : 1.07150 mW
  Ch4 Tx Power          : 1.10180 mW
  Hi Tx Power Alarm Thresh : 3.46730 mW
  Low Tx Power Alarm Thresh : 0.07240 mW
  Ch1 Tx Power Alarm     : None
  Ch2 Tx Power Alarm     : None
  Ch3 Tx Power Alarm     : None
  Ch4 Tx Power Alarm     : None

Rx Power ( 0 to 6.5535 mW)
  Ch1 Rx Power          : 1.13980 mW
  Ch2 Rx Power          : 1.11720 mW
  Ch3 Rx Power          : 1.08800 mW
  Ch4 Rx Power          : 1.16450 mW
  Hi Rx Power Alarm Thresh : 0.33000 mW
  Low Rx Power Alarm Thresh : 1.01830 mW
  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) : 12-05-2016
```

Related Commands

Note

This example is for a QSFP transceiver

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	
Role	admin	
Example	<pre> 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 0010 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 0030 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 0050 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 0070 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 00 05 8d 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 00f0 00 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 0090 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 00b0 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 00d0 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 00f0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 </pre>	

Related Commands


Note

5.2 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.

5.2.1 Configuring Isolated Interfaces

Figure 21: Interface Isolation Example



➤ **To configure isolated interfaces:**

Step 1. Create the VLANs to be used. Run:

```
switch (config) # vlan 2-5
(config vlan 2-5) # exit
```

Step 2. Unlock isolation interface protocol. Run:

```
switch (config) # protocol isolation-group
```

Step 3. Create isolation Group A. Run:

```
switch (config) # isolation-group GroupA
```

Step 4. Assign VLANs 2 and 3 to isolation Group A. Run:

```
(config isolation-group GroupA) # vlan 2-3
(config isolation-group GroupA) # exit
```

Step 5. Create isolation Group B. Run:

```
switch (config) # isolation-group GroupB
```

Step 6. Assign VLANs 4 and 5 to isolation Group B. Run:

```
(config isolation-group GroupB) # vlan 4-5
(config isolation-group GroupB) # exit
```

Step 7. Set Ethernet interfaces 1-3 to access for VLAN 3. Run:

```
(config) # interface ethernet 1/1 switchport access vlan 3
(config) # interface ethernet 1/2 switchport access vlan 3
(config) # interface ethernet 1/3 switchport access vlan 3
```

Step 8. Isolate Ethernet interfaces 1 and 2 and set Ethernet interfaces 3 as privileged. Run:

```
(config) # interface ethernet 1/1-1/2 isolation-group GroupA mode isolated
(config) # interface ethernet 1/3 isolation-group GroupA mode privileged
```

Step 9. Enable isolation Group A. Run:

```
(config) # isolation-group GroupA no shutdown
```

Step 10. Set Ethernet interfaces 4-6 to trunk. Run:

```
(config) # interface ethernet 1/4 switchport mode trunk
(config) # interface ethernet 1/5 switchport mode trunk
(config) # interface ethernet 1/6 switchport mode trunk
```

Step 11. Isolate Ethernet interfaces 4 and 5 and set Ethernet interfaces 6 as privileged. Run:

```
(config) # interface ethernet 1/4-1/5 isolation-group GroupA mode isolated
(config) # interface ethernet 1/6 isolation-group GroupA mode privileged
```

Step 12. Enable isolation Group B. Run:

```
(config) # isolation-group GroupB no shutdown
```

Step 13. Verify configuration. Run:

```
(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
```

5.2.2 Commands

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
Role	admin
Example	switch (config) # protocol isolation-group
Related Commands	
Note	• MLAG must be disabled before enabling interface isolation • When disabled, all configuration is lost

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
Role	admin
Example	switch (config) # isolation-group mygroup
Related Commands	protocol isolation-group
Note	- 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

shutdown

shutdown
no shutdown

Enables isolation group.
The no form of the command disables isolation group.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config Isolation Group
History	3.6.1002
Role	admin
Example	switch (config isolation-group mygroup) # no shutdown
Related Commands	protocol isolation-group isolation-group
Note	Enabling isolation groups fails if there are VLANs with ports both inside and outside the group.

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
Role	admin
Example	switch (config isolation-group mygroup) # vlan 10
Related Commands	protocol isolation-group isolation-group
Note	- 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

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	
Role	admin	
Example	switch (config interface ethernet 1/2) # isolation-group mygroup mode privileged	
Related Commands		
Note	- 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	

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
Role	admin
Example	<pre>switch (config) # show isolation-group mygroup State: Enabled VLANs: 3, 4, 3000 Privileged port: Eth1/25 Isolated ports: Eth1/1, Eth1/2, Eth1/3, Eth1/4, Eth1/5, Eth1/17, Eth1/18, Eth1/19, Eth1/20, Eth1/21, Eth1/27, Eth1/28, Eth1/29, Po600, Po777</pre>
Related Commands	
Note	

5.3 Link Aggregation Group (LAG)

Link Aggregation 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 16 physical ports can be aggregated on a single LAG.

5.3.1 Configuring Static Link Aggregation Group (LAG)

➤ *To configure a static LAG:*

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

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

Step 3. Create a port-channel entity. Run:

```
switch (config) # interface port-channel 1
switch (config interface port-channel 1) #
```

Step 4. Change back to config mode.

```
switch (config interface port-channel 1) # exit
switch (config) #
```

Step 5. Add a physical port to the port-channel. Run:

```
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.

5.3.2 Configuring Link Aggregation Control Protocol (LACP)

➤ *To configure LACP:*

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

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

Step 3. Create a port-channel entity. Run:

```
switch (config) # interface port-channel 1
switch (config interface port-channel 1) #
```

Step 4. Change back to config mode. Run:

```
switch (config interface port-channel 1) # exit  
switch (config) #
```

Step 5. Enable LACP in the switch. Run:

```
switch (config) # lacp  
switch (config) #
```

Step 6. Add a physical port to the port-channel. Run:

```
switch (config interface ethernet 1/4) # channel-group 1 mode active/passive  
switch (config interface ethernet 1/4) #
```

5.3.3 Commands

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	First version
	3.2.1100	Added range support
	3.4.0000	Added note
Role	admin	
Example	<pre>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) #</pre>	
Related Commands		
Note	If a LAG is also an IPL, attempting to delete it without first deleting the IPL is rejected by the management.	

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
Role	admin
Example	switch (config)# lacp switch (config)#
Related Commands	
Note	

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	
Role	admin	
Example	<pre>switch (config)# lacp system-priority 1 switch (config)# show lacp interfaces port-channel Port-channel Module Admin Status is enabled Port-channel System Identifier is 00:02:c9:5c:61:70 LACP System Priority: 3 switch (config)#</pre>	
Related Commands		
Note		

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															
History	3.1.1400															
Role	admin															
Example	<pre>switch (config interface ethernet 1/7)# lacp rate fast switch (config interface ethernet 1/7)# show lacp interfaces ethernet 1/7 Port : 1/7 ----- Port State = Down Channel Group : 1 Pseudo port-channel = Po1 LACP port-priority = 32768 LACP Rate = Slow LACP Activity : Passive LACP Timeout : Short Aggregation State : Aggregation, Defaulted,</pre> <table><thead><tr><th>Port</th><th>State</th><th>LACP Port Priority</th><th>Admin Key</th><th>Oper Key</th><th>Port Number</th><th>Port State</th></tr></thead><tbody><tr><td>1/7</td><td>Down</td><td>128</td><td>1</td><td>1</td><td>0x7</td><td>0x0</td></tr></tbody></table> <pre>switch (config)#</pre>		Port	State	LACP Port Priority	Admin Key	Oper Key	Port Number	Port State	1/7	Down	128	1	1	0x7	0x0
Port	State	LACP Port Priority	Admin Key	Oper Key	Port Number	Port State										
1/7	Down	128	1	1	0x7	0x0										
Related Commands																
Note	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.															

port-channel load-balance ethernet

port-channel load-balance ethernet <method>
no port-channel load-balance ethernet <method>

Configures the port-channel load balancing distribution function method.
 The no form of the command sets the distribution function method to default.

Syntax Description	method	Possible load balance methods: • destination-ip • destination-mac • destination-port • source-destination-ip • source-destination-mac • source-destination-port • source-ip • source-mac • source-port
Default	source-destination-mac	
Configuration Mode	Config	
History	3.1.1400	
Role	admin	
Example	<pre>switch (config)# port-channel load-balance ethernet destination-ip source-port source-mac switch (config)# show interfaces port-channel load-balance destination-ip,source-mac,source-port switch (config)#</pre>	
Related Commands		
Note	Several load balance methods can be configured (refer to the example)	

channel-group

channel-group <1-4096> [mode {on | active | passive}]
no channel-group

Assigns and configures a physical interface to a port channel.
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
Role	admin	
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	
Note	- 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	

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	
Role	admin	
Example	switch (config interface port-channel 10)# lacp-individual enable force	
Related Commands		
Note	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.	

ip address dhcp

ip address dhcp
no ip address dhcp

Enables 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
Role	admin
Example	switch (config interface port channel 10) # ip address dhcp switch (config interface port channel 10) #
Related Commands	interface port-channel show interface port-channel
Note	

show lacp counters

show lacp counters

Displays the LACP PDUs counters.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.1.1400
Role	admin
Example	<pre> switch (config)# show lacp counters LACPDUs Marker Marker Response Port Sent Recv Sent Recv LACPDUs Sent Recv Illegal Unknown ----- Port-channel: 1 ----- 1/7 0 0 0 0 0 0 0 0 switch (config) # switch (config)# </pre>
Related Commands	
Note	

show lacp interfaces ethernet

show lacp interface ethernet <inf>

Displays the LACP interface configuration and status.

Syntax Description	inf	Interface number, for example “1/1”.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.1.1400	
Role	admin	
Example	<pre>switch (config) # show lacp interfaces ethernet 1/4 Port : 1/4 ----- Port State = Down Channel Group : 1 Pseudo port-channel = Po1 LACP port-priority = 128 LACP Rate = Slow LACP Activity : Passive LACP Timeout : Short Aggregation State : Aggregation, Defaulted, Port LACP Port Admin Oper Port Port Port State Priority Key Key Number State ----- 1/4 Down 128 1 1 0x4 0x0 switch (config) #</pre>	
Related Commands		
Note		

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	First version
	3.4.0000	Updated output
Role	admin	

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,

switch (config) #
```

Related Commands

Note

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
Role	admin
Example	switch (config) # show lacp Port-channel Module Admin Status is enabled switch (config) #
Related Commands	
Note	

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	
Role	admin	
Example	switch (config)# show lacp interfaces port-channel 2 system-identifier Priority: 12345 MAC: 00:02:C9:AC:2A:60 switch (config)#	
Related Commands		
Note		

show interfaces port-channel

show interfaces port-channel <port-channel>

Displays port-channel 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
Role	admin	
Example	<pre>switch (config) # show interfaces port-channel 2 Po2 Admin state: Enabled Operational state: Up Description: N/A Mac address: 00:00:00:00:00:00 MTU: 9216 bytes (Maximum packet size 9238 bytes) lacp-individual mode: Enabled Flow-control: receive off send off Actual speed: 2 X 40 Gbps Width reduction mode: disabled Switchport mode: trunk MAC learning mode: Enabled Last clearing of "show interface" counters : Never 60 seconds ingress rate: 2440 bits/sec, 305 bytes/sec, 5 packets/sec 60 seconds egress rate: 2440 bits/sec, 305 bytes/sec, 5 packets/sec Rx 24060 packets 23447 unicast packets 598 multicast packets 15 broadcast packets 1796876 bytes 0 error packets 0 discard packets Tx 23961 packets 23454 unicast packets 496 multicast packets 11 broadcast packets 1805778 bytes 0 error packets 4 discard packets switch (config) #</pre>	

Related Commands

Note

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	
Role	admin	
Example	<pre>switch (config) # show interfaces port-channel 3 counters 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 switch (config) #</pre>	
Related Commands		
Note		

show interfaces port-channel compatibility-parameters

show interfaces port-channel compatibility-parameters

Displays port-channel parameters.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.4000
Role	admin
Example	<pre>switch (config) # show interfaces port-channel compatibility-parameters * Port-mode * Speed * MTU * Flow Control * Access VLAN * Allowed VLAN list * Flowcontrol & PFC * Channel-group mode * CoS 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 switch (config) #</pre>
Related Commands	
Note	

show interfaces port-channel load-balance

show interfaces port-channel load-balance

Displays the type of load-balancing in use for port-channels.

Syntax Description	N/A	N/A
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.4000	
Role	admin	
Example	<pre>switch (config) # show interfaces port-channel load-balance source-destination-mac switch (config) #</pre>	
Related Commands		
Note		

show interfaces port-channel summary


show interfaces port-channel summary

Displays a summary for the port-channel interfaces.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.1.1400 3.4.1100 Updated Example
Role	admin
Example	<pre>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) switch (config) #</pre>
Related Commands	
Note	

5.4 MLAG

Figure 22: Basic MLAG Setup



All nodes in an MLAG must be of the same CPU type (i.e. PPC or x86).



Each switch configuration is independent and it is user responsibility to make sure to configure both switches similarly pertaining MLAG (e.g. MLAG port-channel VLAN membership, static MAC, ACL, etc).

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 currently supported for 2 switches only.



The VIP address must be on the same management IP subnet.

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.



MLAG links currently mandate disabling xSTP control protocol. However, interfaces not part of an MLAG can run any protocol independently.

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 either 10GbE or 40GbE links. In such a case, PFC must be configured on this IPL. [Figure 23, “Basic MLAG Topology,” on page 725](#) illustrates this. The IPL serves the following purposes:

- MLAG protocol control – keepalive messages, MAC sync, MLAG port sync, etc.
- 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. port-channel, Ethernet).

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 is always directed to the MLAG-VIP master node.

When running MLAG with L3, VRRP or MAGP must be deployed. For more information, refer to [Section 6.6, “VRRP,” on page 1258](#) or [Section 6.7, “MAGP,” on page 1273](#) respectively.



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.

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 MLNX-OS software.



When working with MLAG the maximum number of MAC addresses is limited to 47,970. Without it, the number of MAC addresses would be 55,872.



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 features such as Gateway HA, SM HA, and 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*].



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>[<mrag cluster name>:unknown]#
```

This does not reflect the MLAG state, but only the state of the cluster.

5.4.1 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 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 based on UDP packets running via the management plane is maintained and alerts both peers of IPL failure. In such a case of IPL failure, the slave shuts down its interfaces to avoid a split brain scenario and the master becomes a standalone switch.

5.4.2 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.

5.4.3 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.

5.4.4 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.

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.

5.4.5 Upgrading MLAG Pair


Switches in the same MLAG group must have the same MLNX-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`” on page 741) before closing its ports.

For more information on MLAG upgrade, please see Section 4.6.3, “Upgrading MLNX-OS HA Groups,” on page 247.

5.4.6 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 following Mellanox Community post: <https://community.mellanox.com/docs/DOC-2262>.

➤ **To configure L2 MLAG:**

Prerequisites:

Step 1. Enable IP routing. Run:

```
switch (config)# ip routing
```

Step 2. (Recommended) Enable LACP in the switch. Run:

```
switch (config)# lacp
```

Step 3. Enable QoS on the switch to avoid congestion on the IPL port. Run:

```
switch (config)# dcb priority-flow-control enable force
```

Step 4. Enable the MLAG protocol commands. Run:

```
switch (config)# protocol mlag
```

Configuring the IPL:

Step 1. Create a VLAN for the inter-peer link (IPL) to run on. Run:

```
switch (config)# vlan 4000
switch (config vlan 4000)#
```

Step 2. Create a LAG. Run:

```
switch (config)# interface port-channel 1
switch (config interface port-channel 1)#
```

Step 3. Map a physical port to the LAG in active mode (LACP). Run:

```
switch (config)# interface ethernet 1/1 channel-group 1 mode active
```

Step 4. Set this LAG as an IPL. Run:

```
switch (config interface port-channel 1)# ipl 1
```

Step 5. Enable QoS on this specific interface. Run:

```
switch (config interface port-channel 1)# dcb priority-flow-control mode on force
```

Step 6. Create a VLAN interface. Run:

```
switch (config)# interface vlan 4000
switch (config interface vlan 4000)#
```

Step 7. Set an IP address and netmask for the VLAN interface.

On SwitchA, run:

```
switch (config interface vlan 4000)# ip address 10.10.10.1 /30
```

On SwitchB, run:

```
switch (config interface vlan 4000)# ip address 10.10.10.2 /30
```

Step 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 10.10.10.2
```

On SwitchB, run:

```
switch (config interface vlan 4000)# ip 1 peer-address 10.10.10.1
```

Step 9. Configure a virtual IP (VIP) for the MLAG. Run:

On SwitchA, run:

```
switch (config)# mlag-vip my-vip ip 10.10.10.254 /24 //mask may also be 255.255.255.0
```

On SwitchB, run:

```
switch (config)# mlag-vip my-vip
```

Step 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:

Step 1. Create an MLAG interface for the host. Run:

```
switch (config)# interface mlag-port-channel 1
switch (config interface mlag-port-channel 1)#
```

Step 2. Disable STP. Run:

```
switch (config interface mlag-port-channel 1)# spanning-tree port type edge
switch (config interface mlag-port-channel 1)# spanning-tree bpdupfilter enable
```

Step 3. Bind an Ethernet port to the MLAG group. Run:

```
switch (config interface ethernet 1/2)# mlag-channel-group 1 mode on
```

Step 4. Create and enable the MLAG interface. Run:

```
switch (config interface mlag-port-channel 1)# no shutdown
```



STP must be disabled (no spanning-tree) on the MLAG switches when there is at least 1 MLAG port-channel connected to a switch and not to a host.

Enabling MLAG:

Step 1. Enable MLAG. Run:

```
switch [my-vip: master] (config mlag)# no shutdown
```



When running MLAG with L3, VRRP or MAGP must be deployed. For more information, refer to [Section 6.6, “VRRP,” on page 1258](#) or [Section 6.7, “MAGP,” on page 1273](#) respectively.

➤ **To verify MLAG configuration:**

Step 1. Examine MLAG configuration and status. Run:

```
SX2 [mellanox: 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
   Port-Channel Interface State      IP address  IP address
-----
1    Po1         1         Up           10.10.10.1  10.10.10.2

Peers state Summary:
System-id      State  Hostname
-----
F4:52:14:2D:9B:88 Up      <SX2>
F4:52:14:2D:9B:08 Up      SX1
switch [mellanox: master] (config)#

```

Step 2. Examine the MLAG summary table. Run:

```

switch [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
Group
Port-Channel      Type      Local Ports      Peer Ports
(D/P/S/I)          (D/P/S/I)          (D/P/S/I)
-----
1 Mpo2(U)         Static    Eth1/2(P)         Eth1/2(P)

switch (config)#

```

Step 3. Examine the MLAG statistics. Run:

```
switch [my-vip: master] (config)# show mlag statistics
IPL 1:
Rx Heartbeat : 516
Tx Heartbeat : 516
Rx IGMP tunnel : 0
Tx IGMP 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
switch (config)#
```

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.

5.4.7 Commands

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	
Default	no protocol mlag
Configuration Mode	Config
History	3.3.4500
Role	admin
Example	switch (config) # protocol mlag switch (config) #
Related Commands	
Note	- 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.

mlag

mlag

Enters MLAG configuration mode.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.3.4500
Role	admin
Example	switch (config) # mlag switch (config mlag) #
Related Commands	
Note	

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
Role	admin
Example	switch (config mlag) # no shutdown switch (config mlag) #
Related Commands	
Note	This parameter must be similar in all MLAG peers.

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	Integer. Interface number range: 1-1000.
Default	N/A	
Configuration Mode	Config	
History	3.3.4500	
Role	admin	
Example	switch (config) # interface mlag-port-channel 1 switch (config interface mlag-port-channel 1) #	
Related Commands		
Note	• 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.	

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	
Role	admin	
Example	switch (config interface port-channel 1)# ipl 1	
Related Commands		
Note	• 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.	

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	
Role	admin	
Example	switch (config interface vlan 1)# ipl 1 peer-address 10.10.10.10 switch (config interface vlan 1)#	
Related Commands		
Note	• The subnet mask is the same subnet mask of the VLAN interface. • This VLAN interface should be used for IPL only.	

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	
Role	admin	
Example	switch (config mlag) # keep-alive-interval 1 switch (config mlag) #	
Related Commands		
Note	This parameter must be similar in all MLAG peers.	

mlag-channel-group mode

mlag-channel-group <if-number> mode {on | active | passive}
no mlag-channel-group

Binds an Ethernet port to the MLAG LAG.
The no form of the command deletes the binding.

Syntax Description	if-number	Integer. 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	
Role	admin	
Example	switch (config interface ethernet 1/1)# mlag-channel-group 1 mode on switch (config interface ethernet 1/1)#	
Related Commands		
Note		

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
	<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	
Role	admin	
Example	switch (config)# mlag-vip my-mlag-domain ip 10.10.10.254/24 switch (config)#	
Related Commands		
Note	• 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.	

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	
Role	admin	
Example	switch (config mlag) # reload-delay 30 switch (config mlag) #	
Related Commands		
Note	• 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.	

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		
Role	admin		
Example	switch (config mlag) # system-mac 00:00:5E:00:01:5D switch (config mlag) #		
Related Commands			
Note	This parameter must be configured the same in all MLAG peers.		

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	
Role	admin	
Example	switch (config mlag) # upgrade-timeout 60 switch (config mlag) #	
Related Commands		
Note	This parameter must be configured the same in all MLAG peers.	

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	
Role	admin		
Example	<pre>SX2 [mellanox: 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 Port-Channel Interface State IP address IP address ----- 1 Po1 1 Up 10.10.10.1 10.10.10.2 MLAG Members Summary: System-id State Hostname ----- F4:52:14:2D:9B:88 Up <SX2> F4:52:14:2D:9B:08 Up SX1 SX2 [mellanox: master] (config)#</pre>		
Related Commands			
Note			

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
Role	admin
Example	<pre>switch (config)# show mlag-vip MLAG VIP ===== MLAG group name: my-mlag-group MLAG VIP address: 1.1.1.1/30 Active nodes: 2 Hostname VIP-State IP Address ----- SwitchA master 10.10.10.1 SwitchB standby 10.10.10.2 switch (config)#</pre>
Related Commands	
Note	

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
Role	admin
Example	<pre> switch (config)# show interfaces mlag-port-channel 1 Mpo1 Admin state: Enabled Operational state: Down Description: N\A Mac address: 00:00:00:00:00:00 MTU: 1500 bytes (Maximum packet size 1522 bytes) Flow-control: receive off send off Actual speed: 0 Gbps Width reduction mode: disabled Switchport mode: access 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 error packets 0 discard packets switch (config)# </pre>
Related Commands	
Note	

show interfaces mlag-port-channel counters

show interfaces mlag-port-channel <if-number> counters

Displays the extended counters for the interface.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.6.1002
Role	admin
Example	<pre>switch (config)# show interfaces mlag-port-channel 3 counters</pre> <pre> 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 1000000000 error packets 0 discard packets 0 pause packets switch (config)# </pre>
Related Commands	
Note	

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	First version		
	3.4.0000	Added notes and updated example		
	3.4.1100	Updated Example		
Role	admin			
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			
	Group			
	Port-Channel	Type	Local Ports	Peer Ports
	(D/U/P/S)		(D/P/S/I)	(D/P/S/I)

	1 Mpo2(U)	Static	Eth1/2(P)	Eth1/2(P)
	2 Mpo3(U)	Static	Eth1/4(P)	Eth1/8(P)
	3 Mpo4(U)	LACP	Eth1/5(P)	Eth1/5(P)
	switch (config)#			
Related Commands				
Note	• 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 port-channel and in individual state • The “S” flag indicates an interface which is part of a port-channel and in suspended state			

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
Role	admin
Example	<pre>switch (config)# show mlag statistics IPL 1: RX Heartbeat: 439908 TX Heartbeat: 439951 RX IGMP tunnel: 0 TX IGMP tunnel: 1 RX mlag-notification: 0 TX mlag-notification: 12 RX port-notification: 56 TX port-notification: 73 RX FDB sync: 424 TX FDB sync: 778 RX LACP manager: 38 TX LACP manager: 21</pre>
Related Commands	
Note	

5.5 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 take 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 – This mode is Mellanox specific that 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.

5.5.1 Configuring Access Mode and Assigning Port VLAN ID (PVID)

➤ *To configure Access mode and assign PVID to interfaces:*

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

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

Step 3. Create a VLAN. Run:

```
switch (config) # vlan 6
switch (config vlan 6) #
```

Step 4. Change back to config mode. Run:

```
switch (config vlan 6) # exit
switch (config) #
```

Step 5. Enter the interface context. Run:

```
switch (config) # interface ethernet 1/36
switch (config interface ethernet 1/36) #
```

Step 6. From within the interface context, configure the interface mode to Access. Run:

```
switch (config interface ethernet 1/36) # switchport mode access
switch (config interface ethernet 1/36) #
```

Step 7. From within the interface context, configure the Access VLAN membership. Run:

```
switch (config interface ethernet 1/36) # switchport access vlan 6
switch (config interface ethernet 1/36) #
```

Step 8. Change back to config mode. Run:

```
switch (config interface ethernet 1/36) # exit
switch (config) #
```

5.5.2 Configuring Hybrid Mode and Assigning Port VLAN ID (PVID)

➤ *To configure Hybrid mode and assign PVID to interfaces:*

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

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

Step 3. Create a VLAN. Run:

```
switch (config) # vlan 6
switch (config vlan 6) #
```

Step 4. Change back to config mode. Run:

```
switch (config vlan 6) # exit
switch (config) #
```

Step 5. Enter the interface context. Run:

```
switch (config) # interface ethernet 1/36
switch (config interface ethernet 1/36) #
```

Step 6. From within the interface context, configure the interface mode to Access. Run:

```
switch (config interface ethernet 1/36) # switchport mode hybrid
switch (config interface ethernet 1/36) #
```

Step 7. From within the interface context, configure the Access VLAN membership. Run:

```
switch (config interface ethernet 1/36) # switchport access vlan 6
switch (config interface ethernet 1/36) #
```

Step 8. Change to config mode again. Run:

```
switch (config interface ethernet 1/36) # exit
switch (config) #
```

5.5.3 Configuring Trunk Mode VLAN Membership

➤ *To configure Trunk mode VLAN membership:*

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

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

Step 3. Create a VLAN. Run:

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

Step 4. Change back to config mode. Run:

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

Step 5. Enter the interface context. Run:

```
switch [standalone: master] (config) # interface ethernet 1/35
switch [standalone: master] (config interface ethernet 1/35) #
```

Step 6. From within the interface context, configure the interface mode to Trunk. Run:

```
switch [standalone: master] (config interface ethernet 1/35) # switchport mode trunk
switch [standalone: master] (config interface ethernet 1/35) #
```

5.5.4 Configuring Hybrid Mode VLAN Membership

➤ *To configure Hybrid mode VLAN membership:*

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

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

Step 3. Create a VLAN. Run:

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

Step 4. Change back to config mode. Run:

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

Step 5. Enter the interface context. Run:

```
switch (config) # interface ethernet 1/35
switch (config interface ethernet 1/35) #
```

Step 6. From within the interface context, configure the interface mode to Hybrid. Run:

```
switch (config interface ethernet 1/35) # switchport mode hybrid
switch (config interface ethernet 1/35) #
```

Step 7. From within the interface context, configure the allowed VLAN membership. Run:

```
switch (config interface ethernet 1/35) # switchport hybrid allowed-vlan add 10
switch (config interface ethernet 1/35) #
```

Step 8. Change to config mode again. Run:

```
switch (config interface ethernet 1/35) # exit
switch (config) #
```

5.5.5 Commands

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	1-4094.
	vlan-range	Any range of VLANs.
Default	VLAN 1 is enabled by default.	
Configuration Mode	Config	
History	3.1.1400	
Role	admin	
Example	<pre>switch (config) # vlan 10 switch (config vlan 10) # show vlan</pre>	
	<pre>VLAN Name Ports ---- - 1 default Eth1/2, Eth1/3, Eth1/4/1, Eth1/4/2 ... 10 switch (config vlan 10) #</pre>	
Related Commands	<pre>show vlan switchport mode switchport [trunk hybrid] allowed-vlan</pre>	
Note	Interfaces are not added automatically to VLAN unless configured with trunk or hybrid mode with “all” option turned on.	

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																									
Role	admin																									
Example	<pre>switch (config) # vlan 10 switch (config vlan 10) # name my-vlan-name switch (config vlan 10) # show vlan</pre> <table> <thead> <tr> <th>VLAN</th><th>Name</th><th>Ports</th></tr> </thead> <tbody> <tr> <td>1</td><td>default</td><td>Eth1/2, Eth1/3, Eth1/4/1, Eth1/4/2, Eth1/5,</td></tr> <tr> <td>15,</td><td></td><td>Eth1/6, Eth1/7, Eth1/8, Eth1/9, Eth1/10, Eth1/11, Eth1/12, Eth1/13, Eth1/14, Eth1/15,</td></tr> <tr> <td>20,</td><td></td><td>Eth1/16, Eth1/17, Eth1/18, Eth1/19, Eth1/20,</td></tr> <tr> <td>25,</td><td></td><td>Eth1/21, Eth1/22, Eth1/23, Eth1/24, Eth1/25,</td></tr> <tr> <td>30,</td><td></td><td>Eth1/26, Eth1/27, Eth1/28, Eth1/29, Eth1/30,</td></tr> <tr> <td>35,</td><td></td><td>Eth1/31, Eth1/32, Eth1/33, Eth1/34, Eth1/35,</td></tr> <tr> <td>10</td><td>my-vlan-name</td><td>Eth1/36, Po34, Po4096</td></tr> </tbody> </table>		VLAN	Name	Ports	1	default	Eth1/2, Eth1/3, Eth1/4/1, Eth1/4/2, Eth1/5,	15,		Eth1/6, Eth1/7, Eth1/8, Eth1/9, Eth1/10, Eth1/11, Eth1/12, Eth1/13, Eth1/14, Eth1/15,	20,		Eth1/16, Eth1/17, Eth1/18, Eth1/19, Eth1/20,	25,		Eth1/21, Eth1/22, Eth1/23, Eth1/24, Eth1/25,	30,		Eth1/26, Eth1/27, Eth1/28, Eth1/29, Eth1/30,	35,		Eth1/31, Eth1/32, Eth1/33, Eth1/34, Eth1/35,	10	my-vlan-name	Eth1/36, Po34, Po4096
VLAN	Name	Ports																								
1	default	Eth1/2, Eth1/3, Eth1/4/1, Eth1/4/2, Eth1/5,																								
15,		Eth1/6, Eth1/7, Eth1/8, Eth1/9, Eth1/10, Eth1/11, Eth1/12, Eth1/13, Eth1/14, Eth1/15,																								
20,		Eth1/16, Eth1/17, Eth1/18, Eth1/19, Eth1/20,																								
25,		Eth1/21, Eth1/22, Eth1/23, Eth1/24, Eth1/25,																								
30,		Eth1/26, Eth1/27, Eth1/28, Eth1/29, Eth1/30,																								
35,		Eth1/31, Eth1/32, Eth1/33, Eth1/34, Eth1/35,																								
10	my-vlan-name	Eth1/36, Po34, Po4096																								
Related Commands	<pre>show vlan switchport mode switchport [trunk hybrid] allowed-vlan</pre>																									
Note	Name can not be added to a range of VLANs.																									

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										
Role	admin										
Example	<pre>switch (config vlan 10) # show vlan</pre> <table> <thead> <tr> <th>VLAN</th><th>Name</th><th>Ports</th></tr> </thead> <tbody> <tr> <td>1</td><td>default</td><td>Eth1/2, Eth1/3, Eth1/4/1, Eth1/4/2 ...</td></tr> <tr> <td>10</td><td>my-vlan-name</td><td></td></tr> </tbody> </table>		VLAN	Name	Ports	1	default	Eth1/2, Eth1/3, Eth1/4/1, Eth1/4/2 ...	10	my-vlan-name	
VLAN	Name	Ports									
1	default	Eth1/2, Eth1/3, Eth1/4/1, Eth1/4/2 ...									
10	my-vlan-name										
Related Commands	<pre>show vlan switchport mode switchport [trunk hybrid] allowed-vlan vlan</pre>										
Note											

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 MLAG port-channel configuration mode
	3.4.3000	Added dot1q-tunnel parameter
Role	admin	
Example	<pre>switch (config) # interface ethernet 1/7 switch (config interface ethernet 1/7) # switchport mode access switch (config interface ethernet 1/7) # show interfaces switchport Interface Mode Access vlan Allowed vlans ----- ----- ----- ----- Eth1/2 access 1 Eth1/3 access 1 Eth1/4/1 access 1 Eth1/4/2 access 1 Eth1/5 access 1 Eth1/6 access 1 Po34 access 1 Po4096 access 1 switch (config interface ethernet 1/7) #</pre>	

Related Commands

- show vlan
- show interfaces switchport
- switchport access vlan
- switchport [trunk | hybrid] allowed-vlan
- switchport dot1q-tunnel qos-mode
- vlan

Note

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 switch (config interface ethernet 1/1) #	
Related Commands	show vlan show interfaces switchport switchport access vlan switchport [trunk hybrid] allowed-vlan vlan	
Note		

switchport access

switchport access vlan <vlan-id>
no switchport access vlan

Sets the port access VLAN.
The no form of the command sets the port access VLAN to 1.

Syntax Description	vlan-id1-4094.																																									
Default	1																																									
Configuration Mode	Config Interface Ethernet Config Interface Port Channel Config Interface MLAG Port Channel																																									
History	3.1.1400	First version																																								
	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 MLAG port-channel configuration mode																																								
Role	admin																																									
Example	<pre>switch (config) # interface ethernet 1/7 switch (config interface ethernet 1/7) # switchport access vlan 10 switch (config interface ethernet 1/7) # show interfaces switchport</pre> <table><thead><tr><th>Interface</th><th>Mode</th><th>Access vlan</th><th>Allowed vlans</th></tr></thead><tbody><tr><td>Eth1/2</td><td>access</td><td>1</td><td></td></tr><tr><td>Eth1/3</td><td>access</td><td>1</td><td></td></tr><tr><td>Eth1/4/1</td><td>access</td><td>1</td><td></td></tr><tr><td>Eth1/4/2</td><td>access</td><td>1</td><td></td></tr><tr><td>Eth1/5</td><td>access</td><td>1</td><td></td></tr><tr><td>Eth1/6</td><td>access</td><td>1</td><td></td></tr><tr><td>Eth1/7</td><td>access</td><td>10</td><td></td></tr><tr><td>....</td><td></td><td></td><td></td></tr><tr><td>Po4096</td><td>access</td><td>1</td><td></td></tr></tbody></table> <pre>switch (config interface ethernet 1/7) #</pre>		Interface	Mode	Access vlan	Allowed vlans	Eth1/2	access	1		Eth1/3	access	1		Eth1/4/1	access	1		Eth1/4/2	access	1		Eth1/5	access	1		Eth1/6	access	1		Eth1/7	access	10						Po4096	access	1	
Interface	Mode	Access vlan	Allowed vlans																																							
Eth1/2	access	1																																								
Eth1/3	access	1																																								
Eth1/4/1	access	1																																								
Eth1/4/2	access	1																																								
Eth1/5	access	1																																								
Eth1/6	access	1																																								
Eth1/7	access	10																																								
....																																										
Po4096	access	1																																								
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 applicable to configure on the port, depends on the switchport mode of the port.																																									

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 except 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	
Role	admin	
Example	<pre> switch (config) # interface ethernet 1/7 switch (config interface ethernet 1/7) # switchport hybrid allowed-vlan all switch (config interface ethernet 1/7) #show interfaces switchport Interface Mode Access vlan Allowed vlans ----- ----- ----- ----- Eth1/2 access 1 Eth1/3 access 1 Eth1/4/1 access 1 Eth1/4/2 access 1 Eth1/5 access 1 Eth1/6 access 1 Eth1/7 hybrid 1 1, 10 Po34 access 1 Po4096 access 1 switch (config interface ethernet 1/7) # </pre>	

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.

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																																									
Role	admin																																									
Example	<pre>switch (config) # interface ethernet 1/7 switch (config interface ethernet 1/7) # switchport voice vlan 10 switch (config interface ethernet 1/7) # show interfaces switchport</pre> <table><thead><tr><th>Interface</th><th>Mode</th><th>Access vlan</th><th>Allowed vlans</th></tr></thead><tbody><tr><td>Eth1/2</td><td>access</td><td>1</td><td></td></tr><tr><td>Eth1/3</td><td>access</td><td>1</td><td></td></tr><tr><td>Eth1/4/1</td><td>access</td><td>1</td><td></td></tr><tr><td>Eth1/4/2</td><td>access</td><td>1</td><td></td></tr><tr><td>Eth1/5</td><td>access</td><td>1</td><td></td></tr><tr><td>Eth1/6</td><td>access</td><td>1</td><td></td></tr><tr><td>Eth1/7</td><td>access</td><td>10</td><td></td></tr><tr><td>....</td><td></td><td></td><td></td></tr><tr><td>Po4096</td><td>access</td><td>1</td><td></td></tr></tbody></table> <pre>switch (config interface ethernet 1/7) #</pre>		Interface	Mode	Access vlan	Allowed vlans	Eth1/2	access	1		Eth1/3	access	1		Eth1/4/1	access	1		Eth1/4/2	access	1		Eth1/5	access	1		Eth1/6	access	1		Eth1/7	access	10						Po4096	access	1	
Interface	Mode	Access vlan	Allowed vlans																																							
Eth1/2	access	1																																								
Eth1/3	access	1																																								
Eth1/4/1	access	1																																								
Eth1/4/2	access	1																																								
Eth1/5	access	1																																								
Eth1/6	access	1																																								
Eth1/7	access	10																																								
....																																										
Po4096	access	1																																								
Related Commands	lldp med-tlv-select show vlan show interfaces switchport switchport mode switchport [trunk hybrid] allowed-vlan vlan																																									
Note																																										

show interface switchport

show interface switchport

Displays all interface switch port configurations.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.1.1400
Role	admin
Example	<pre>switch (config) #show interfaces switchport Interface Mode Access vlan Allowed vlans ----- ----- ----- ----- Eth1/2 access 1 Eth1/3 access 1 Eth1/4/1 access 1 Eth1/4/2 access 1 Eth1/5 access 1 Eth1/6 access 1 Eth1/7 hybrid 1 1, 10 Po34 access 1 Po4096 access 1 switch (config)#</pre>
Related Commands	<pre>show vlan switchport access vlan switchport mode vlan</pre>
Note	

5.6 Voice VLAN


This feature 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.

MLNX-OS® allow the user to configure the VLAN for voice traffic. In [Figure 24](#), 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.

Figure 24: Tagging Voice Packets with a Different VLAN ID



5.6.1 Configuring Voice VLAN

➤ *To configure LLDP-MED TLV, run:*

```
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:**

```
switch (config) # show lldp interfaces
TLV flags:
PD: port-description, SN: sys-name, SD: sys-description, SC: sys-capabilities, MA: man-
agement-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 interfaces ethernet 1/4
TLV flags:
PD: port-description, SN: sys-name, SD: sys-description, SC: sys-capabilities, MA: man-
agement-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 interfaces 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:**

Step 1. Create a VLAN. Run:

```
switch (config) # vlan 200
switch (config vlan 200) # exit
switch (config) #
```

Step 2. Set the interface mode to be hybrid. Run:

```
switch (config) # interface ethernet 1/4 switchport mode hybrid
switch (config) # interface ethernet 1/4 switchport hybrid allowed-vlan 200
```

Step 3. Assign the VLAN to the interface. Run:

```
switch (config) # interface ethernet 1/4 switchport voice vlan 200
```

Step 4. (Optional) Change the PVID of the port so that untagged packets go to a different VLAN than the default. Run:

```
switch (config)# vlan 300
switch (config vlan 300)# exit
switch (config)# interface ethernet 1/4 switchport access vlan 300
```

Step 5. Verify the configuration. Run:

```
switch (config)# show interfaces 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 interfaces ethernet 1/4
TLV flags:
PD: port-description, SN: sys-name, SD: sys-description, SC: sys-capabilities, MA: man-
agement-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 interfaces 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:**

Step 1. Remove the voice VLAN from the interface. Run:

```
switch (config)# no interface ethernet 1/4 switchport voice vlan
```

Step 2. Disable the MED TLV from the interface. Run:

```
switch (config)# interface ethernet 1/4 lldp med-tlv-select none
```

5.6.2 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.

5.7 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.

5.7.1 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.

5.7.2 Configuring QinQ

➤ To configure QinQ:

Step 1. Create the C-VLAN. Run:

```
switch (config) # vlan 200
switch (config vlan 200) # exit
```

Step 2. Enter the configuration mode of an Ethernet, LAG, or MLAG interface. Run:

```
switch (config) # interface port-channel 100
```

Step 3. Change the switchport mode of the interface to enable QinQ. Run:

```
switch (config interface port-channel 100) # switchport mode dot1q-tunnel
```

Step 4. Change its port VLAN ID (PVID). This configures the S-VLAN. Run:

```
switch (config interface port-channel 100) # switchport access vlan 200
```

Step 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
switch (config interface port-channel 100) #
```

Step 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
switch (config interface port-channel 100) #
```


5.7.3 Commands

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	
Role	admin	
Example	switch (config interface ethernet 1/1) # switchport dot1q-tunnel qos-mode uniform switch (config interface ethernet 1/1) #	
Related Commands	show vlan show interfaces switchport switchport access vlan switchport [trunk hybrid] allowed-vlan vlan	
Note		

5.8 MAC Address Table

5.8.1 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:**

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

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

Step 3. Run the command “mac-address-table static unicast <destination mac address> vlan <vlan identifier(1-4094)> interface ethernet <slot>/<port>”.


```
switch (config) # mac-address-table static 00:11:22:33:44:55 vlan 1 interface ethernet
1/1
```

5.8.2 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) 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.

Figure 25: MAC Learning Disable Example Case



5.8.3 Commands

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	
Role	admin	
Example	<pre>switch (config) # mac-address-table aging-time 50 switch (config) # show mac-address-table aging-time Mac Address Aging Time: 50 switch (config) #</pre>	
Related Commands	<pre>show mac-address-table show mac-address-table aging time</pre>	
Note		

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	The interface number (i.e. 1/1, 3).
Default	No static MAC addresses available in default.	
Configuration Mode	Config	
History	3.1.0600	
Role	admin	
Example	<pre>switch (config) # mac-address-table static aa:aa:aa:aa:aa:aa vlan 1 interface ethernet 1/7 switch (config) # show mac-address-table Switch ethernet-default Vlan Mac Address Type Interface ---- - 1 aa:aa:aa:aa:aa:aa static Eth1/7 Number of unicast: 1 Number of multicast: 0 switch (config) #</pre>	
Related Commands	<pre>show mac-address-table mac-address-table aging time</pre>	
Note	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.	

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
Role	admin
Example	switch (config interface ethernet 1/1) # mac-learning disable
Related Commands	
Note	• 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.

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
Role	admin
Example	switch (config) # clear mac-address-table dynamic switch (config) #
Related Commands	mac-address-table aging-time mac-address-table static show mac-address-table
Note	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.

show mac-address-table

show mac-address-table [**address** <mac-address> | **interface ethernet** <if-number> | **vlan** [<vlan> | **range** <range>] | **unicast** | **multicast**]

Displays the static and dynamic unicast and multicast MAC addresses for the switch. Various of filter options available.

Syntax Description	mac-address	Filter the table to a specific MAC address.
	if-number	Filter the table to a specific interface.
	vlan	Filter the table to a specific VLAN number (1-4094).
	range	Filter the table to a range of VLANs.
	unicast	Filter the table to a unicast addresses only.
	multicast	Filter the table to a multicast addresses only.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.1.0600	
	3.3.4500	Updated Example
Role	admin	
Example	<pre>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 Number of multicast: 0 switch (config) #</pre>	
Related Commands	mac-address-table static clear mac-address-table	
Note		

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
Role	admin
Example	<pre>switch (config) # mac-address-table aging-time 300 switch (config) # show mac-address-table aging-time</pre> Mac Address Aging Time: 300 <pre>switch (config) #</pre>
Related Commands	<pre>mac-address-table aging-time mac-address-table static clear mac-address-table</pre>
Note	MAC addresses learned on the mgmt0 is not shown by this command.

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
Role	admin
Example	switch (config) # show mac-address-table summary Number of unicast: 4 Number of multicast: 0
Related Commands	mac-address-table static clear mac-address-table
Note	

5.9 Spanning Tree

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 Mellanox 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.

5.9.1 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 <1-2000000000>
```

5.9.2 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 to a switch. If it does not receive any BPDUs, it is moved to discarding state.

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}
```

5.9.3 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}
```



Configuring BPDU filtering on a port connected to a switch can cause bridging loops because the port filters any BPDU it receives and goes to forwarding state.

5.9.4 BPDU Guard

BPDU guard is a security feature which, when enabled, shuts down the port 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 port type <type> bpduguard
```

5.9.5 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” port type.

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
```

5.9.6 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
```

5.9.7 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.


For MSTP network design over Mellanox L2 VMS, please refer to [Mellanox Virtual Modular Switch Reference Guide](#).

5.9.8 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.

Figure 26 exhibits a typical RPVST network configuration to get a better utilization on the inter-switch trunk ports.

Figure 26: RPVST Network Config



5.9.8.1 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 in MLNX-OS. Enforcement of this limitation is done through the maximum number of VLANs allowed in the system. On x86 switch systems the limitation is 128 VLANs, and on PPC systems it ranges from 13-18 VLANs depending on the switch system. The more ports the switch system has the less VLANs it can support.

Table 57 - Supported VLANs by RPVST per Switch System

Switch System Model	Number of Supported VLANs
x86 systems	128
SX1012	17
SX1016	13
SX1024	13
SX1035	13
SX1036	13


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.

5.9.8.2 RPVST and RSTP Interoperability

Figure 27: RPVST and RSTP Cluster



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.

5.9.9 Commands

spanning-tree

spanning-tree
no spanning-tree

Globally enables the spanning tree feature.
The no form disables the spanning tree feature.

Syntax Description	N/A
Default	Spanning tree is enabled.
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # no spanning-tree switch (config) #
Related Commands	show spanning-tree
Note	

spanning-tree mode

spanning-tree mode {rst | mst | rpvst}
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	Multiple spanning tree.
	rst	Rapid spanning tree.
	rpvst	Rapid per-VLAN spanning tree.
Default	rst	
Configuration Mode	Config	
History	3.3.4150	
Role	admin	
Example	switch (config)# spanning-tree mode mst	
Related Commands		
Note	- On x86 switch systems, the number of VLANs supported by RPVST are 128 - On PPC switch systems, the number of VLANs supported by RPVST are between 13-18	

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]

Sets the 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	
Role	admin	
Example	switch (config) # spanning-tree forward-time switch (config) #	
Related Commands	show spanning-tree	
Note	The following formula applies on the spanning tree timers: $2 * (\text{ForwardTime} - 1) \geq \text{MaxAgeTime} \geq 2 * (\text{Hello Time} + 1)$	

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
Role	admin	
Example	switch (config) # spanning-tree port type edge default switch (config) #	
Related Commands	show spanning-tree	
Note		

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. Its value must be in steps of 4096, starting from 0. Only the following values are applicable: 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	
Role	admin	
Example	switch (config) # spanning-tree priority 4096 switch (config) #	
Related Commands	show spanning-tree	
Note		

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. The possible values are: 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 MLAG port-channel configuration mode
Role	admin	
Example	switch (config) # interface ethernet 1/1 switch (config interface ethernet 1/1) # spanning-tree port-priority 16 switch (config interface ethernet 1/1) #	
Related Commands	show spanning-tree	
Note		

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. Value range is 0-2000000000.
Default	The default cost is derived from the speed. 1Gbps 20000 10Gbps 2000 40Gbps 500 56Gbps 357	
Configuration Mode	Config Interface Ethernet Config Interface Port Channel Config Interface MLAG Port Channel	
History	3.1.0000	
	3.3.4500	Added MLAG port-channel configuration mode
Role	admin	
Example	switch (config) # interface ethernet 1/1 switch (config interface ethernet 1/1) # spanning-tree cost 1000 switch (config interface ethernet 1/1) #	
Related Commands	show spanning-tree	
Note	- 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 (UIP 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	

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 MLAG port-channel configuration mode
Role	admin	
Example	switch (config) # interface ethernet 1/1 switch (config interface ethernet 1/1) # spanning-tree port type edge switch (config interface ethernet 1/1) #	
Related Commands	show spanning-tree	
Note		

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 loop-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 MLAG port-channel configuration mode
Role	admin	
Example	switch (config) # interface ethernet 1/1 switch (config interface ethernet 1/1) # spanning-tree guard root switch (config interface ethernet 1/1) #	
Related Commands	show spanning-tree	
Note		

spanning-tree bpdupfilter

spanning-tree bpdupfilter {disable | enable}
no spanning-tree bpdupfilter

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	
Role	admin	
Example	switch (config) # interface ethernet 1/1 switch (config interface ethernet 1/1) # spanning-tree bpdupfilter enable	
Related Commands	show spanning-tree	
Note	This command can be used when the switch is connected to hosts.	

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
Role	admin
Example	switch (config) # clear panning-tree counters switch (config) #
Related Commands	show spanning tree
Note	

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. The range is 6-40.
Default	20	
Configuration Mode	Config	
History	3.3.4150	
Role	admin	
Example	switch (config)# spanning-tree mst max-hops 20 switch (config)#	
Related Commands		
Note	- 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	

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 is 1-64.
	priority	MST instance port priority. Possible values are: 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	
Role	admin	
Example	switch (config)# spanning-tree mst 1 priority 32768 switch (config)#	
Related Commands		
Note	• 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	

spanning-tree mst vlan

spanning-tree mst <mst-instance> vlan <vlan-range>
no spanning-tree mst <mst-instance> vlan <vlan-range>

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 is 1-64.
	vlan <vlan-range>	A single VLAN or a a range of VLANs. The format is <vlan> or <from-vlan>-<to-vlan>.
Default	N/A	
Configuration Mode	Config	
History	3.3.4150	
Role	admin	
Example	switch (config)# spanning-tree mst 1 vlan 10-20 switch (config)#	
Related Commands		
Note	This command is available when global STP mode is set to MST	

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	The MST revision number. Range is 0-65535.
Default	0	
Configuration Mode	Config	
History	3.3.4150	
Role	admin	
Example	switch (config)# spanning-tree mst revision 1 switch (config)#	
Related Commands		
Note	- 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	

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	
Role	admin	
Example	switch (config)# spanning-tree mst name my-mst switch (config)#	
Related Commands		
Note	- 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	

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. Possible range is 1-64.
	role	Values: “primary” or “secondary”.
Default	primary	
Configuration Mode	Config	
History	3.3.4150	
Role	admin	
Example	switch (config)# spanning-tree mst name my-mst switch (config)#	
Related Commands		
Note	• 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.	

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 is 0-4094.
	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	
Role	admin	
Example	switch (config interface ethernet 1/1)# spanning-tree mst 1 port-priority 32768 switch (config interface port-channel 1)# spanning-tree mst 1 port-priority 32768	
Related Commands		
Note	This command is available when global STP mode is set to MST.	

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 is 1-64.
	cost-value	MST instance port cost. Range is 0-200000000.
Default	2000 for 10Gb/s, 500 for 40Gb/s, 20000 for 1Gb/s, 357 for 56Gb/s	
Configuration Mode	Config Interface Port Channel	
History	3.3.4150	
Role	admin	
Example	<pre>switch (config interface ethernet 1/1)# spanning-tree mst 1 cost 4000 switch (config interface port-channel 1)# spanning-tree mst 1 cost 4000 switch (config)#</pre>	
Related Commands		
Note	This command is available when global STP mode is set to MST.	

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	
Role	admin	
Example	switch (config) # spanning-tree vlan 10 forward-time 15	
Related Commands	show spanning-tree	
Note	- The following formula applies on the spanning tree timers: $2 * (\text{ForwardTime} - 1) \geq \text{MaxAgeTime} \geq 2 * (\text{Hello Time} + 1)$ - This command is available when global STP mode is set to RPVST	

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	secs	Parameter range: 1-2 seconds.
Default	2 seconds	
Configuration Mode	Config	
History	3.4.1100	
Role	admin	
Example	switch (config) # spanning-tree vlan 10 hello-time 2	
Related Commands	show spanning-tree	
Note	- The following formula applies on the spanning tree timers: $2 * (\text{ForwardTime} - 1) \geq \text{MaxAgeTime} \geq 2 * (\text{Hello Time} + 1)$ - This command is available when global STP mode is set to RPVST	

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	Parameter range: 6-40 seconds.
Default	20 seconds	
Configuration Mode	Config	
History	3.4.1100	
Role	admin	
Example	switch (config) # spanning-tree vlan 10 max-age 20	
Related Commands	show spanning-tree	
Note	- The following formula applies on the spanning tree timers: $2 * (\text{ForwardTime} - 1) \geq \text{MaxAgeTime} \geq 2 * (\text{Hello Time} + 1)$ - This command is available when global STP mode is set to RPVST	

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	priority	Possible values are: 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	
Role	admin	
Example	switch (config) # spanning-tree vlan 10 priority 32768	
Related Commands	show spanning-tree	
Note	- The following formula applies on the spanning tree timers: $2 * (\text{ForwardTime} - 1) \geq \text{MaxAgeTime} \geq 2 * (\text{Hello Time} + 1)$ - This command is available when global STP mode is set to RPVST	

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
Role	admin
Example	<pre> switch (config) # show spanning-tree Switch ethernet-default Spanning tree protocol is enabled rst Spanning tree force version:2 Root ID Priority 32768 Address 00:02:c9:7a:e9:40 Cost 1000 Port Eth1/32 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec Bridge ID Priority 32768 Address 00:02:c9:96:c6:d0 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec L - Loop Inconsistent R - Root Inconsistent G - BPDU Guard Inconsistent Interface Role Sts Cost Prio Type ---- Eth1/9 Designated Forwarding 500 128 normal Eth1/22 Designated Discarding(R) 500 128 normal Eth1/32 Root Forwarding 500 128 normal Eth1/39 Disabled Discarding(G) 2000 128 normal switch (config) # </pre>
Related Commands	<pre> clear spanning-tree counters spanning-tree </pre>
Note	

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
Role	admin
Example	<pre>switch (config) # show spanning-tree detail Switch ethernet-default Spanning tree protocol is enabled Bridge is executing the rst compatible Spanning Tree Protocol Bridge Identifier has priority 32768, address 00:02:c9:96:c6:d0 Configured hello time 2, max age 20, forward delay 15 Current root has priority 32768, address 00:02:c9:7a:e9:40 Root port is Eth1/32(Ethernet1/32),cost of root path is 1000 Number of topology changes 21,last change occurred 00:00:03 ago Timers: hold 6 hello 2, max age 20, forward delay 15 default port type: normal, default bpdu filter: disabled, default bpdu guard: disabled switch (config) #</pre>
Related Commands	<pre>clear spanning-tree counters spanning-tree</pre>
Note	

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	
Role	admin	
Example	<pre>switch (config) # show spanning-tree interface ethernet 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 switch (config) #</pre>	
Related Commands	<pre>clear spanning-tree counters spanning-tree</pre>	
Note		

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	
Role	admin	
Example	<pre>switch (config) # show spanning-tree mst MST0 vlans mapped: 1-1023,1025-2047,2049-3071,3073-4094 Interface Role Sts Cost Prio Type ---- Eth1/9 Designated Forwarding 500 128.9 point-to-point Eth1/10 Designated Forwarding 500 128.10 point-to-point Eth1/11 Back Up Discarding 500 128.22 point-to-point switch (config) #</pre>	
Related Commands	<pre>clear spanning-tree counters spanning-tree</pre>	
Note		

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																																																
Role	admin																																																
Example	switch (config) # show spanning-tree root <table><thead><tr><th>Instance</th><th>Priority</th><th>MAC addr</th><th>Root Cost</th><th>Hello Time</th><th>Max Age</th><th>FWD Dly</th><th>Root Port</th></tr><tr><th>-----</th><th>-----</th><th>-----</th><th>-----</th><th>-----</th><th>-----</th><th>-----</th><th>-----</th></tr></thead><tbody><tr><td>MST0</td><td>32768</td><td>00:02:c9:71:ed:40</td><td>500</td><td>2</td><td>20</td><td>15</td><td>Eth1/20</td></tr><tr><td>MST1</td><td>32768</td><td>00:02:c9:71:f0:c0</td><td>0</td><td>2</td><td>20</td><td>15</td><td>-</td></tr><tr><td>MST2</td><td>0</td><td>00:02:c9:71:f0:c0</td><td>0</td><td>2</td><td>20</td><td>15</td><td>-</td></tr><tr><td>MST3</td><td>32768</td><td>00:02:c9:71:f0:c0</td><td>0</td><td>2</td><td>20</td><td>15</td><td>-</td></tr></tbody></table> switch (config) #	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	-
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																																																
Note																																																	

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 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	
Role	admin	
Example	<pre> switch (config) # show spanning-tree vlan 10 Switch ethernet-default Spanning tree protocol is enabled rpvst Spanning tree force version:2 Vlan 10 Root ID Priority 10 Address 00:02:c9:96:c6:d0 This bridge is the root Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec Bridge ID Priority 10 Address 00:02:c9:96:c6:d0 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec L - Loop Inconsistent Interface Role Sts Cost Prio Type ---- - Mpo21 Designated Forwarding 500 128 normal Mpo20 Back Up Discarding 500 128 normal switch (config) # </pre>	

Related Commands	clear spanning-tree counters spanning-tree

Note

5.10 OpenFlow

MLNX-OS supports OpenFlow 1.0 (on SwitchX®) and 1.3 (on Spectrum™). 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.

Mellanox 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.

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.

5.10.1 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, Ether-Type, 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.

Upkeeping 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



The flow table on SwitchX® supports up to 1000 flows.

5.10.2 Configuring OpenFlow

➤ *To run OpenFlow on a switch:*

Step 1. Unlock the OpenFlow CLI commands. Run:

```
switch (config) # protocol openflow
```

Step 2. Configure interfaces to be managed by OpenFlow. Run:

```
switch (config) # interface ethernet 1/1-1/4 openflow mode hybrid
```

Step 3. Configure the OpenFlow controller IP and TCP port. Run:

```
switch (config) # openflow controller-ip 10.209.0.205 tcp-port 6633
```



Spectrum based systems do not support a different controller port other than the default (6633).

Step 4. (Optional) Verify the OpenFlow configuration. Run:

```
switch (config) # show openflow
OpenFlow version: OF VERSION 1.0
Table size: 1000, 0 in use
Active controller ip: 10.209.0.205 port: 6633
Connection status: HANDSHAKE_COMPLETE (CONNECTED)
Forward-to-controller: ospf lldp arp-unicast arp-broadcast (all)
Enabled ports:  Eth1/1      Eth1/2      Eth1/3      Eth1/4
switch (config) #
```



To be able to configure the switch using the controller, you should see the following line in the output:
Connection status must be: HANDSHAKE_COMPLETE (CONNECTED).

5.10.3 Commands

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
Role	admin
Example	switch (config) # protocol openflow switch (config) #
Related Commands	
Note	

openflow description (SwitchX)

openflow description <string>

Sets the OpenFlow description.

Syntax Description	string	Free string.
Default	N/A	
Configuration Mode	Config	
History	3.3.4302	
	3.6.1002	Updated Note
Role	admin	
Example	<pre>switch (config) # openflow description OF-switch-104 switch (config) # show openflow detail OpenFlow version: OF VERSION 1.0 Table size: 1000, 0 in use Active controller ip: 10.209.1.39 port: 6633 Connection status: HANDSHAKE_COMPLETE (CONNECTED) Forward-to-controller: ospf lldp arp-unicast arp-broadcast (all) Enabled ports: Eth1/10 Eth1/11 Eth1/13 Eth1/19 Echo period: 10 sec Keep alive period: 30 sec Messages in (last session): 86290 Messages out (last session): 47984 Disconnect count: 0 Openflow description: OF-switch-104 Datapath ID: 00:00:00:02:c9:a8:e3:50 Not supporting buffering Not supporting emergency flows Not supporting port statistics Not supporting IP reassemble Supporting spanning tree Not supporting queue statistics switch (config) #</pre>	
Related Commands		
Note	Not supported on Spectrum based switch systems	

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
History	3.3.4200 3.6.2100 Updated Note section
Role	admin
Example	switch (config interface ethernet 1/1)# openflow mode hybrid switch (config interface ethernet 1/1)#
Related Commands	
Note	On Spectrum based systems, it is possible to run “interface port-channel <port number> openflow mode hybrid”

controller-ip

openflow controller-ip <ip-address> [tcp-port <tcp-port>]
no openflow controller-ip [tcp-port <tcp-port>]

Sets the OpenFlow controller's IP & TCP port.
 The no form of the command sets the parameter to its default.

Syntax Description	ip-address	The IPv4 address of the OpenFlow controller
	tcp-port	Sets the TCP port number of the OpenFlow controller
Default	TCP port 6633	
Configuration Mode	Config OpenFlow	
History	3.3.4200	
Role	admin	
Example	switch (config openflow) # controller-ip 10.10.10.10 tcp-port 6633	
Related Commands		
Note	This command is only supported on SwitchX based switch systems	

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 is 0x0000-0xFFFF in hexa.
Default	0x0000	
Configuration Mode	Config OpenFlow	
History	3.3.4200	
Role	admin	
Example	switch (config openflow) # datapath-id 0x1234 switch (config openflow) #	
Related Commands		
Note		

forward-to-controller

forward-to-controller {[ospf] [lldp] [arp-unicast] [arp-broadcast] all | none}

Forwards the selected traffic types to the controller from all the ports on which OpenFlow enabled.

Syntax Description	ospf	Forwards OSPF traffic to the controller
	lldp	Forwards LLDP traffic to the controller
	arp-unicast	Forwards ARP-unicast traffic to the controller
	arp-broadcast	Forwards ARP-broadcast traffic to the controller
	all	Forwards all traffic types to the controller
	none	Forwards no traffic to the controller
Default	None	
Configuration Mode	Config OpenFlow	
History	3.3.4200	
Role	admin	
Example	<pre>switch (config openflow) # forward-to-controller all switch (config openflow) #</pre>	
Related Commands		
Note	This command is only supported on SwitchX based switch systems	

show openflow detail

show openflow detail

Displays detailed information about the OpenFlow protocol.

Syntax Description	N/A
Default	None
Configuration Mode	Any Command Mode
History	3.3.4200
	3.6.1002 Updated Example
Role	admin
Example	<pre>switch (config) # show openflow detail Echo period: 0 sec Keep alive period: 0 sec Messages in (last session): 0 Messages out (last session): 0 Disconnect count: 0 Openflow description: Datapath ID: 02:10:e4:52:14:5d:76:70 Not supporting buffering Not supporting emergency flows Not supporting port statistics Not supporting IP reassemble Supporting spanning tree Not supporting queue statistics</pre>
Related Commands	
Note	This command is only supported on SwitchX based switch systems

show openflow flows

show openflow flows

Displays information about the OpenFlow flows.

Syntax Description	N/A
Default	None
Configuration Mode	Any Command Mode
History	3.3.4302 3.6.1002 Updated Example
Role	admin
Example	<pre>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,d_l_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,d_l_vlan=10 actions=output:123 cookie=0x0, duration=468.013s, table=0, n_packets=0, n_bytes=0, send_flow_rem priority=8,d_l_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,d_l_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</pre>
Related Commands	
Note	

show openflow statistics

show openflow statistics

Displays information about the OpenFlow flows.

Syntax Description	N/A
Default	None
Configuration Mode	Any Command Mode
History	3.3.4302
	3.6.1002 Updated Example
Role	admin
Example	switch (config) # show openflow statistics
Related Commands	
Note	This command is only supported on SwitchX based switch systems

show openflow tables

show openflow tables

Displays information about the OpenFlow tables (size, type, etc.).

Syntax Description	N/A
Default	None
Configuration Mode	Any Command Mode
History	3.3.4200
	3.6.1002 Added Example
Role	admin
Example	<pre>switch (config) # show openflow tables Table id: 0 Maximum table size: 1000, 0 in use Key: 12 tuple ACL Supported actions: Modify VID, Mofify PCP, Strip VID</pre>
Related Commands	
Note	This command is only supported on SwitchX based switch systems

show openflow

show openflow [detail | tables | flows <id>]

Displays general information about the OpenFlow protocol configuration.

Syntax Description	detail	Displays detailed information about the OpenFlow protocol.
	tables	Displays information about the OpenFlow tables (size, type, etc.).
	flows <id>	Displays specific flows inside the OpenFlow tables. ID may be a range (e.g. 1-10).
	statistics	Displays OpenFlow statistics.
Default	None	
Configuration Mode	Any Command Mode	
History	3.3.4200	
	3.3.4302	Removed flow-id parameter Added “flows” and “statistics” parameters
Role	admin	
Example	switch (config) # show openflow flows	
Related Commands		
Note	This command is only supported on SwitchX based switch systems	

5.11 OVS VTEP

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 that 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.

5.11.1 Configuring OVS VTEP

➤ *To enable VTEP:*

Step 1. Configure jumbo frames for NVE ports. Run:

```
switch (config)# interface ethernet 1/1-1/4 mtu 9216 force
```

Step 2. Configure jumbo frames for ESXi facing ports. Run:

```
switch (config)# interface ethernet 1/17 mtu 9216 force
```

Step 3. Create VLAN for all VXLAN traffic. Run:

```
switch (config)# vlan 3
```

Step 4. Configure ESXi interfaces with VXLAN VLAN. Run:

```
switch (config)# interface ethernet 1/17 switchport access vlan 3
```

Step 5. Enable IP routing. Run:

```
switch (config)# ip routing vrf default
```

Step 6. Create loopback interface and configure an IP address for it (this is the switch VTEP IP). Run:

```
switch (config)# interface loopback 1
switch (config interface loopback 1)# ip address 1.2.3.4 255.255.255.255
```

Step 7. Configure interface on the VXLAN VLAN and configure an IP address for it (this should be the default gateway of the VTEP subnet on the ESXi servers). 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
```

Step 8. Enable NVE protocol. Run:

```
switch (config)# protocol nve
```

Step 9. Configure interface NVE. Run:

```
switch (config)# interface nve 1
```

Step 10. Configure the source of the NVE interface to be the loopback create above. Run:

```
switch (config)# interface nve 1 vxlan source interface loopback 1
```

Step 11. Start OVSDB server. Run:

```
switch (config)# ovs ovsdb server
```

Step 12. Configure the OVSDB manager to an IP address of an NSX controller. Run:

```
switch (config)# ovs ovsdb manager remote ssl ip address 10.130.250.5
```

Step 13. Configure switch 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
```

5.11.2 Commands

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 data-base.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.6.3004
Role	admin
Example	switch (config) # protocol nve
Related Commands	
Note	This command is only supported on Spectrum™ based switch systems

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	
Role	admin	
Example	switch (config) # interface nve 1 switch (config interface nve 1) #	
Related Commands	protocol nve	
Note	This command is only supported on Spectrum™ based switch systems	

shutdown

shutdown
no shutdown

Disables VXLAN tunnel.
The no form of the command enables VXLAN tunnel.

Syntax Description	N/A
Default	Enabled
Configuration Mode	Config Interface NVE
History	3.6.3004
Role	admin
Example	switch (config interface nve 1) # shutdown
Related Commands	interface nve protocol nve
Note	This command is only supported on Spectrum™ based switch systems

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 Valid range: 0-31
Default	N/A	
Configuration Mode	Config Interface NVE	
History	3.6.3004	
Role	admin	
Example	switch (config interface nve 1) # vxlan source interface loopback 14	
Related Commands	interface nve protocol nve	
Note	The configured loopback interface becomes the VXLAN tunnel endpoint (VTEP) This command is only supported on Spectrum™ based switch systems	

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	Not in NVE mode	
Configuration Mode	Config Interface Ethernet	
History	3.6.3004	
Role	admin	
Example	switch (config interface ethernet 1/1) # nve mode only	
Related Commands	interface ethernet	
Note	This command is only supported on Spectrum™ based switch systems	

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
Role	admin
Example	switch (config interface nve 1) # clear nve counters
Related Commands	interface nve protocol nve
Note	This command is only supported on Spectrum™ based switch systems

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																				
Role	admin																				
Example	<pre>switch (config) # show interface nve</pre> <table><thead><tr><th>Remote Manager IP Address</th><th>Port</th><th>Connection Type</th></tr><tr><th>-----</th><th>----</th><th>-----</th></tr></thead><tbody><tr><td>2.2.2.2</td><td>200</td><td>tcp</td></tr></tbody></table> <pre>NVE member interfaces: Eth1/2, Eth1/7</pre> <pre>Interface NVE 1 status: Admin state: up Source interface: loopback 1</pre> <table><tbody><tr><td>17971</td><td>encapsulated (Tx) NVE packets</td></tr><tr><td>0</td><td>decapsulated (Rx) NVE packets</td></tr><tr><td>0</td><td>dropped NVE-encapsulated packets</td></tr><tr><td>0</td><td>NVE-encapsulated packets with errors</td></tr></tbody></table>				Remote Manager IP Address	Port	Connection Type	-----	----	-----	2.2.2.2	200	tcp	17971	encapsulated (Tx) NVE packets	0	decapsulated (Rx) NVE packets	0	dropped NVE-encapsulated packets	0	NVE-encapsulated packets with errors
Remote Manager IP Address	Port	Connection Type																			
-----	----	-----																			
2.2.2.2	200	tcp																			
17971	encapsulated (Tx) NVE packets																				
0	decapsulated (Rx) NVE packets																				
0	dropped NVE-encapsulated packets																				
0	NVE-encapsulated packets with errors																				
Related Commands																					
Note	This command is only supported on Spectrum™ based switch systems																				

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	
Role	admin	
Example	<pre>switch (config) # show interface nve 1 counters 18330 encapsulated (Tx) NVE packets 0 decapsulated (Rx) NVE packets 0 dropped NVE-encapsulated packets 0 NVE-encapsulated packets with errors</pre>	
Related Commands		
Note	This command is only supported on Spectrum™ based switch systems	

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			
Role	admin			
Example	switch (config) # show interface nve 1 flood			
	NVE Interface	Logical Switch	VNI ID	Flood IP Address
	-----	-----	-----	-----
	1	ls7777	7777	1.2.3.5
Related Commands				
Note	This command is only supported on Spectrum™ based switch systems			

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 MAC address table of NVE interface with specified VNI				
Default	N/A					
Configuration Mode	Any Command Mode					
History	3.6.3004					
Role	admin					
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						
Note	This command is only supported on Spectrum™ based switch systems					

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 MAC addresses on the bridge with the given VNI				
Default	N/A					
Configuration Mode	Any Command Mode					
History	3.6.3004					
Role	admin					
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						
Note	This command is only supported on Spectrum™ based switch systems					

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 only MAC addresses on the bridge with the given VNI				
Default	N/A					
Configuration Mode	Any Command Mode					
History	3.6.3004					
Role	admin					
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						
Note	This command is only supported on Spectrum™ based switch systems					

show interfaces nve mac-address-table remote configured unicast

show interfaces nve <nve-id> mac-address-table remote configured unicast [vni <vni-id>]

Displays only remotely-configured unicast addresses.

Syntax Description	nve-id	NVE ID range: 1-64				
	vni	Displays only MAC addresses on the bridge with the given VNI				
Default	N/A					
Configuration Mode	Any Command Mode					
History	3.6.3004					
Role	admin					
Example	switch (config) # show interface nve 1 mac-address-table remote configured unicast					
	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						
Note	This command is only supported on Spectrum™ based switch systems					

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 peers on specific VNI		
Default	N/A			
Configuration Mode	Any Command Mode			
History	3.6.3004			
Role	admin			
Example	switch (config) # show interface nve 1 peers			
	NVE Interface	Logical Switch	VNI ID	Peer IP Address
	-----	-----	-----	-----
	1	ls7777	7777	1.2.3.5
Related Commands				
Note	This command is only supported on Spectrum™ based switch systems			

ovs ovssdb server

ovs ovssdb server
no ovs ovssdb server

Runs OVSSDB-server process and unhides OVS commands.
 The no form of the command deactivates OVSSDB-server process and hides OVS commands.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.6.3004
Role	admin
Example	switch (config) # ovs ovssdb server
Related Commands	
Note	This command is only supported on Spectrum™ based switch systems

ovs ovssdb manager remote

```
ovs ovssdb manager remote {tcp | ssl} ip-address <ip-address> port <tcp-port>
no ovs ovssdb manager remote {tcp | ssl} ip-address <ip-address> port <tcp-port>
```

Configures OVSSDB 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 OVSSDB 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	
Role	admin	
Example	switch (config) # ovs ovssdb manager remote tcp ip-address 10.10.10.10 port 20	
Related Commands	ovs ovssdb server	
Note	This command is only supported on Spectrum™ based switch systems	

ovs ovssdb server listen

ovs ovssdb server listen {tcp | ssl} port <tcp-port> local ip-address <ip-address>
no ovs ovssdb server listen {tcp | ssl} port <tcp-port> local ip-address <ip-address>

Configures OVSSDB to listen at a given port of an interface with a given (local) IP address.

The no form of the command disconnects OVSSDB 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	
Role	admin	
Example	switch (config) # ovs ovssdb server listen tcp port 20 local ip-address 20.20.20.20	
Related Commands	ovs ovssdb server	
Note	This command is only supported on Spectrum™ based switch systems	

5.12 IGMP Snooping



While IGMPv3 is supported on SwitchX®, the source is not considered. So a “join” to a group from a specific source (S,G) is treated as a join to the group from all sources (*,G).

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 data-base 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.

5.12.1 Configuring IGMP Snooping

You can configure IGMP snooping to establish multicast group memberships.

➤ **To configure IGMP snooping:**

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

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

Step 3. Enable IGMP snooping globally. Run:

```
switch (config) # ip igmp snooping
switch (config) #
```

Step 4. Enable IGMP snooping on a VLAN. Run:

```
switch (config) # vlan 2
switch (config vlan 2) # ip igmp snooping
```

5.12.2 Defining a Multicast Router Port on a VLAN

You can define a Multicast Router (MRouter) port on a VLAN in one of the following methods:

➤ **To change the interface switchport to trunk:**

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

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

Step 3. Enable IGMP snooping globally. Run:

```
switch (config) # ip igmp snooping
switch (config) #
```

Step 4. 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
```

Step 5. Change back to config mode. Run:

```
switch (config interface ethernet 1/1) # exit
switch (config) #
```

Step 6. Define the MRouter port on the VLAN. Run:

```
switch (config) # vlan 2
switch (config vlan 2) # ip igmp snooping mrouter interface ethernet 1/1
switch (config vlan 2) #
```

➤ **To change the interface switchport to hybrid:**

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

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

Step 3. Enable IGMP snooping globally. Run:

```
switch (config) # ip igmp snooping
switch (config) #
```

Step 4. Create a VLAN. Run:

```
switch (config) # vlan 200
switch (config vlan 200) #
```

Step 5. Change back to config mode. Run:

```
switch (config vlan 200) # exit
switch (config) #
```

Step 6. Change the interface switchport mode of the port (the interface is member of VLAN 1 by default). Run:

```
switch (config) # interface ethernet 1/36
switch (config interface ethernet 1/36) # switchport mode hybrid
```

Step 7. Attach the VLAN to the port's interface. Run:

```
switch (config interface ethernet 1/36) # switchport mode hybrid allowed-vlan 200
switch (config interface ethernet 1/36) #
```

Step 8. Change to config mode again. Run:

```
switch (config interface ethernet 1/36) # exit
switch (config) #
```

Step 9. Define the MRouter port on the VLAN. Run:

```
switch (config) # vlan 200
switch (config vlan 200) # ip igmp mrouter interface ethernet 1/36
switch (config vlan 200) #
```

5.12.3 IGMP Snooping Querier

IGMP Snooping Querier compliments 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:**

Step 1. Enable the IGMP snooping on the switch. Run:

```
switch (config) # ip igmp snooping
```

Step 2. Enable the IGMP snooping querier on a specific VLAN. Run:

```
switch (config) # vlan 10
switch (config vlan 10)# ip igmp snooping querier
```

Step 3. Set the query interval time. Run:

```
switch (config vlan 10)# ip igmp snooping querier query-interval 25
```

Step 4. (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

switch (config vlan 10)#
```


5.12.4 Commands

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
Role	admin
Example	switch (config) # ip igmp snooping switch (config) # vlan 10 switch (config vlan 10) # ip igmp snooping
Related Commands	show ip igmp snooping
Note	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.

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>}

no ip igmp snooping {last-member-query-interval | proxy reporting | mrouter-timeout | report-suppression-interval}

Configures IGMP global parameters.

The no form of the command resets the IGMP global 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 will be 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. The default value is 125 seconds.
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. The default value is 5 seconds.
Default	last-member-query-interval – 1 second proxy reporting is disabled mrouter-timeout – 125 port-purge-timeout – 260 seconds report-suppression-interval – 5 seconds
Configuration Mode	Config
History	3.1.1400
Role	admin
Example	switch (config) # ip igmp snooping report-suppression-interval 3

Related Commands	ip igmp snooping (admin) show ip igmp snooping

Note

ip igmp snooping clear counters

ip igmp snooping clear 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	
Role	admin	
Example	switch (config) # ip igmp snooping clear counters vlan 2	
Related Commands		
Note		

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	Normal-leave is enabled.
Configuration Mode	Config Interface Ethernet Config Interface Port Channel Config Interface MLAG Port Channel
History	3.1.1400 3.3.4500 Added MLAG port-channel configuration mode
Role	admin
Example	switch (config interface ethernet 1/1) # ip igmp snooping fast-leave
Related Commands	show ip igmp snooping interfaces
Note	

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> type - ethernet or port-channel.
Default	No static mrouters are configured.
Configuration Mode	Config VLAN
History	3.1.1400
Role	admin
Example	switch (config)# vlan 1 switch (config vlan 1) # ip igmp snooping mrouter interface ethernet 1/1
Related Commands	show ip igmp snooping mrouter
Note	The multicast router port can be created only if IGMP snooping is enabled both globally and on the VLAN.

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
Role	admin	
Example	switch (config vlan 1) # ip igmp snooping static-group 230.0.0.1 interface ethernet 1/1	
Related Commands	show ip igmp snooping groups	
Note	If the deleted interface is the last port, it deletes the entire multicast group.	

ip igmp snooping unregistered multicast

ip igmp snooping unregistered multicast <options>
no ip igmp snooping unregistered multicast

Sets the behavior of the snooping switch for unregistered multicast traffic.
The no form of the command sets it default.

Syntax Description	options • flood • forward-to-mrouter-ports
Default	flood
Configuration Mode	Config
History	3.2.0500
Role	admin
Example	switch (config) # ip igmp snooping unregistered multicast flood
Related Commands	show ip igmp snooping
Note	

ip igmp snooping version

ip igmp snooping version {2 | 3}

Configures the default operating version to be used for newly created IGMP snooping instances.

Syntax Description	2	Enables IGMPv2
	3	Enables IGMPv3
Default	3	
Configuration Mode	Config Config VLAN	
History	3.6.1002	
	3.6.2100	Updated default
Role	admin	
Example	switch (config vlan 2)# ip igmp snooping version 3	
Related Commands		
Note		

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
Role	admin
Example	switch (config vlan 1)# ip igmp snooping querier switch (config vlan 1)#
Related Commands	igmp snooping querier query-interval show ip igmp snooping querier
Note	

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	
Role	admin	
Example	switch (config vlan 1)# igmp snooping querier query-interval 20 switch (config vlan 1)#	
Related Commands	igmp snooping querier query-interval show ip igmp snooping querier	
Note		

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 output
Role	admin
Example	<pre>switch (config) # show ip igmp snooping IGMP snooping global configuration: IGMP snooping globally enabled IGMP default version for new VLAN is V3 IGMP snooping operationally enabled Proxy-reporting globally disabled Last member query interval is 1 seconds Mrouter timeout is 125 seconds Port purge timeout is 260 seconds Report suppression interval is 5 seconds IGMP snooping unregistered multicast: flood switch (config) #</pre>
Related Commands	
Note	

show ip igmp snooping groups

show ip igmp snooping groups [vlan <vlan ID> [group <group IP>]]

Displays per VLAN the list of multicast groups attached (static or dynamic allocated) per port.

Syntax Description	N/A				
	<table> <tr> <td>vlan</td><td>VLAN ID</td></tr> <tr> <td>group</td><td>Multicast group IP address</td></tr> </table>	vlan	VLAN ID	group	Multicast group IP address
vlan	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				
Role	admin				
Example	<pre>switch (config) # show ip igmp snooping groups Vlan ID Source St/Dyn Ports ----- 1 12.10.10.1 Dyn Eth1/2 1 12.11.11.2 St Eth1/1 Total Num of Dynamic Group Addresses 1 Total Num of Static Group Addresses 1 switch (config) # show ip igmp snooping groups vlan 1 group 224.5.5.5 Snooping group information for VLAN 1 and group 224.5.5.5 Filter Mode: INCLUDE Include sources: 1.2.3.4 V1/V2 Receiver Ports: None V3 Receiver Ports: Port Number: Eth1/1 Include sources: 1.2.3.4 Exclude sources: None</pre>				
Related Commands					
Note					

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
Role	admin
Example	<pre>switch (config) # show ip igmp snooping interfaces interface leave-mode ----- 1/1 Normal 1/2 Normal 1/3 Normal 1/4 Fast ... switch (config) #</pre>
Related Commands	
Note	

show ip igmp snooping membership

show ip igmp snooping membership [vlan <VID> [group <group IP>]]

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	3.6.2100	
Role	admin	
Example	<pre>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</pre>	
Related Commands		
Note		

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
Role	admin
Example	<pre>switch (config) # show ip igmp snooping mrouter Vlan Ports ----- 1 Eth1/1(static) switch (config) #</pre>
Related Commands	
Note	

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
Role	admin	
Example	<pre>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</pre>	
Related Commands		
Note		

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	
Role	admin	
Example	<pre>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</pre>	
Related Commands		
Note		

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
Role	admin
Example	<pre>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</pre>
Related Commands	
Note	

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	
Role	admin	
Example	<pre>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 switch (config) #</pre>	
Related Commands		
Note		

5.13 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.

5.13.1 Configuring LLDP

➤ *To configure the LLDP on the switch:*

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

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

Step 3. Enable LLDP globally on the switch. Run:

```
switch (config) # lldp
switch (config) #
```

Step 4. Enable LLDP per interface. Run:

```
switch (config interface ethernet 1/1) # lldp receive
switch (config interface ethernet 1/1) # lldp transmit
```

Step 5. Show LLDP local information. Run:

```
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
```

Step 6. Show LLDP remote information. Run:

```
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: ethernet 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
```

5.13.2 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) compliments the DCB implementation by offering a dynamic protocol that communicates DCB attributes between peering endpoint.

MLNX-OS 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, however, is not enabled by default).

For more information, please refer to the Mellanox Community at:

<https://community.mellanox.com/docs/DOC-2485>.

5.13.3 Commands

lldp

lldp
no lldp

Enables LLDP globally.
The no form of the command disables the LLDP.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config
History	3.2.0300
Role	admin
Example	switch (config)# lldp switch (config)#
Related Commands	show lldp local
Note	

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	
Role	admin	
Example	switch (config)# lldp reinit 10 switch (config)#	
Related Commands	show lldp timers	
Note		

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	
Role	admin	
Example	switch (config)# lldp timer 10 switch (config)#	
Related Commands	show lldp timers	
Note		

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	
Role	admin	
Example	switch (config)# lldp tx-delay 10 switch (config)#	
Related Commands	show lldp timers	
Note	The recommended value for the tx-delay is set by the following formula: $1 \leq \text{lldp tx-delay} \leq (0.25 * \text{lldp timer})$	

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	
Role	admin	
Example	switch (config)# lldp tx-hold-multiplier 10 switch (config)#	
Related Commands	show lldp timers	
Note	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.	

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	
Role	admin	
Example	switch (config interface ethernet 1/1)# lldp receive switch (config interface ethernet 1/1)#	
Related Commands	show lldp interface	
Note	The LLDP is disabled by default (globally)	

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	Initial revision
	3.3.0000	Added “none” parameter
	3.3.4302	Added “dcbx” parameter
	3.3.4402	Added “dcbx-cee” parameter
Role	admin	
Example	switch (config interface ethernet 1/1)# lldp tlv-select port-description sys-name switch (config interface ethernet 1/1)#	
Related Commands	show lldp interface	
Note		

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	
Role	admin	
Example	<pre>switch (config interface ethernet 1/1)# lldp med-tlv-select all switch (config interface ethernet 1/1)#</pre>	
Related Commands	show lldp interface	
Note		

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	
	3.4.0008	
Role	admin	
Example	<pre>switch (config-if)# dcb application-priority ethertype 0x8906 switch (config-if)#</pre>	
Related Commands	show lldp interface	
Note		

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
Role	admin
Example	<pre>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: SX1036 Supported capabilities: B,R Supported capabilities enabled: B switch (config)#</pre>
Related Commands	
Note	

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	First version
	3.3.4200	Updated Example
	3.6.1002	Updated Example
Role	admin	
Example	<pre>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</pre>	
Related Commands		
Note		

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				
Role	admin				
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	arc-switch108	
	Eth1/11	e4:1d:2d:44:65:00	Eth1/11	arc-switch108	
	Eth1/12	e4:1d:2d:44:65:00	Eth1/12	arc-switch108	
	Eth1/13	e4:1d:2d:44:65:00	Eth1/13	arc-switch108	
	Eth1/14	e4:1d:2d:44:65:00	Eth1/14	arc-switch108	
	Eth1/15	e4:1d:2d:44:65:00	Eth1/15	arc-switch108	
	Eth1/16	e4:1d:2d:44:65:00	Eth1/16	arc-switch108	
	Eth1/17	e4:1d:2d:44:65:00	Eth1/17	arc-switch108	
	Eth1/18	e4:1d:2d:44:65:00	Eth1/18	arc-switch108	
	Eth1/19	e4:1d:2d:44:65:00	Eth1/19	arc-switch108	
	Eth1/20	e4:1d:2d:44:65:00	Eth1/20	arc-switch108	
	Eth1/21	e4:1d:2d:44:65:00	Eth1/21	arc-switch108	
	Eth1/22	e4:1d:2d:44:65:00	Eth1/22	arc-switch108	
	Eth1/23	e4:1d:2d:44:65:00	Eth1/23	arc-switch108	
	Eth1/24	e4:1d:2d:44:65:00	Eth1/24	arc-switch108	
	Eth1/25	e4:1d:2d:44:65:00	Eth1/25	arc-switch108	
	Eth1/26	e4:1d:2d:44:65:00	Eth1/26	arc-switch108	
	Eth1/31	e4:1d:2d:44:65:00	Eth1/31	arc-switch108	
	Eth1/32	e4:1d:2d:44:65:00	Eth1/32	arc-switch108	
	Related Commands				
	Note				

show lldp statistics [interface ethernet <inf>]

show lldp statistics [interface ethernet <inf>]

Displays LLDP interface statistics.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.2.0300
Role	admin
Example	<pre>switch (config)# show lldp statistics ethernet 1/1 Interface Frames In In TLVs TLVs Ageout Out Discarded Errors Total Discarded Unrecognize Frames ----- Eth 1/1 0 0 10 0 0 0 0 switch (config)#</pre>
Related Commands	
Note	

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
Role	admin
Example	<pre>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 switch (config)#</pre>
Related Commands	
Note	

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
Role	admin
Example	<pre>switch (config)# show lldp timers msg-tx-interval:30 tx-delay:2 tx-hold:4 tx-reinit-delay:2 switch (config)#</pre>
Related Commands	
Note	

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												
Role	admin												
Example	<pre>switch (config)# show dcb application-priority</pre> Application priority configuration <table><thead><tr><th>Selector</th><th>Protocol</th><th>Priority</th></tr></thead><tbody><tr><td colspan="3">-----</td></tr><tr><td>Ethertype</td><td>0x8906</td><td>3</td></tr><tr><td>Ethertype</td><td>0x8914</td><td>3</td></tr></tbody></table> <pre>switch (config)#</pre>	Selector	Protocol	Priority	-----			Ethertype	0x8906	3	Ethertype	0x8914	3
Selector	Protocol	Priority											

Ethertype	0x8906	3											
Ethertype	0x8914	3											
Related Commands													
Note													

5.14 Quality of Service (QoS)

5.14.1 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.

5.14.1.1 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)
 - Else, based on packet DSCP field for IP packet
 - Else, based on port default setting for non-IP
- Trust both
 - Else, based on packet DSCP for IP packet
 - Else, based on packet PCP,DEI for VLAN tagged packets
 - Else, based on the port default setting

Table 58 and figure summarize the classification rules.

Table 58 - 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.

5.14.1.2 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.

5.14.1.3 Default QoS Configuration

Table 59 - 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

Table 59 - Default QoS Configuration

Parameter	Range	Configuration
Switch-priority to IEEE priority	7	7

5.14.2 QoS Rewrite

Spectrum™ based switch systems 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.

5.14.2.1 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.

5.14.2.2 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 affect after deriving switch-priority from the DSCP field.

5.14.2.3 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.

5.14.2.4 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

5.14.3 Queuing and Scheduling (ETS) for SwitchX

Enhanced Transmission Selection (ETS) 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 that 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

The default mapping of priority to traffic classes (per interface) is as follows:

- Priority 0,1 mapped to TC 0
- Priority 2,3 mapped to TC 1
- Priority 4,5 mapped to TC 2
- Priority 6,7 mapped to TC 3



TC0 and TC3 are lossy TCs, while TC1 and TC2 can be lossless as well as lossy. It is possible but not recommended to map PFC enabled priorities (lossless traffic) to those TC0 or TC3.

ETS is enabled by default (scheduling is WRR).

➤ **To set the scheduling mode to Strict Priority:**

Step 1. Run the command `dcb ets disable`.

```
switch (config) # no dcb ets enable
```

➤ **To configure the WRR bandwidth percentage:**

Step 1. Make sure ETS feature is enabled. Run:

```
switch (config) # dcb ets enable
```

Step 2. Choose the WRR bandwidth rate and distribution.

By default the WRR distribution function is equal 25% per TC. Changing the WRR bandwidth rate will cause a change in the distribution function, for example if you wish to

schedule more traffic on TC-0, TC-1, TC-2 while reducing the amount of traffic sent on TC-3, run the command `dcb ets tc bandwidth`.

```
switch (config) # dcb ets tc bandwidth 30 30 30 10
# show dcb ets
```

ETS enabled

TC	Bandwidth
0	30%
1	30%
2	30%
3	10%

Number of Traffic Class: 4

```
switch (config) #
```



Traffic class priorities are <0-3>, where 0 is the lowest and 3 is the highest.



The sum of all traffic class bandwidth value (in percentage) should be 100, otherwise the command fails.

Step 3. Run the command `show dcb ets` to verify the configuration.

```
switch (config) # show dcb ets
ETS enabled
```

TC	Bandwidth
0	30%
1	30%
2	10%
3	30%

Number of Traffic Class: 4

```
switch (config) #
```

5.14.4 Queuing and Scheduling (ETS) for Spectrum

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.

5.14.4.1 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.

5.14.4.2 Multicast Aware Traffic Class Mapping

Spectrum™ supports a mode of MC aware TC mapping if the mapping to the TCs is based also on the whether the packet is unicast or multicast. So, packets of the same switch-priority can be mapped to two different TCs, based on their traffic type. With MC aware mode enabled, MC traffic is mapped into 8 MC TCs in parallel to 8 unicast TCs. Unicast TC has strict priority over its parallel multicast TC.

5.14.4.3 Traffic Shapers

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) [bytes/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.

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.

5.14.4.4 Default Shaper Configuration

Table 60 - 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

Table 60 - Default Shaper Configuration

Parameter	Range	Configuration
Switch-priority to TC	4	4
Switch-priority to TC	5	5
Switch-priority to TC	6	6
Switch-priority to TC	7	7
MC-aware TC mapping	All ports	True
Shaping	All ports	No max/min shaping configured

5.14.5 RED and ECN




Supported only on Spectrum™ based switch systems.

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.

Spectrum™ based systems support relative RED/ECN on TC queues. This feature 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 [Figure 28](#):

Figure 28: RED/ECN Drop Profiles



- 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”.



Spectrum™ based systems support RED/ECN only for unicast traffic classes.

5.14.6 Commands

5.14.6.1 QoS Classification

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	
Role	admin	
Example	switch (config interface ethernet 1/1) # vlan default priority 0	
Related Commands	N/A	
Notes	This command is only supported on Spectrum™ based switch systems	

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
Role	admin
Example	switch (config interface ethernet 1/1) # vlan default dei 0
Related Commands	N/A
Notes	This command is only supported on Spectrum™ based switch systems

qos trust

qos trust [port | L2 | L3 | both]
no qos trust [port | L2 | L3 | both]

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
Role	admin
Example	switch (config interface ethernet 1/1) # qos trust L2
Related Commands	N/A
Notes	This command is only supported on Spectrum™ based switch systems

qos default switch-priority

qos default switch-priority [<switch-priority>]
no qos default switch-priority [<switch-priority>]

Configures default switch-priority for 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	
Role	admin	
Example	switch (config interface ethernet 1/1) # qos default switch-priority 0	
Related Commands	qos trust	
Notes	This command is only supported on Spectrum™ based switch systems	

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	
Role	admin	
Example	switch (config interface ethernet 1/1) # qos map dscp 45	
Related Commands	qos trust	
Notes	This command is only supported on Spectrum™ based switch systems	

show qos

show qos [interface <type> <number>]

Displays QoS information.

Syntax	Description	N/A
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	Any Command Mode	
History	3.6.1002	
Role	admin	

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: preserve (router is disabled)
DSCP rewrite: disabled

PCP,DEI to switch-priority mapping:
PCP,DEI  switch-priority
-----
0,0      0
1,0      1
2,0      2
...
6,1      6
7,1      7

DSCP to switch-priority mapping:
DSCP     switch-priority
-----
0        0
1        0
2        0
...
62       7
63       7

PCP,DEI rewrite mapping (switch-priority to PCP,DEI):
switch-priority  PCP,DEI
-----
0                0,0
1                1,0
2                2,0
...

DSCP rewrite mapping (switch-priority to DSCP):
switch-priority  DSCP
-----
0                0
1                8
2               16
...
```

Related Commands N/A

Notes This command is only supported on Spectrum™ based switch systems

show interfaces ethernet counters tc

show interfaces ethernet <slot/port> counters tc <priority>

Displays traffic 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	
Role	admin	
Example	switch (config) # show interfaces ethernet 1/1 counters tc 3	
	TC 3	
	0	packets
	0	bytes
	0	queue depth
	0	unicast no buffer discard
	0	WRED discard
Related Commands		
Note	This command is only supported on Spectrum™ based switch systems	

show interfaces ethernet counters pg

show interfaces ethernet <slot/port> counters pg <priority>

Displays port 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	
Role	admin	
Example	switch (config) # show interfaces ethernet 1/1 counters pg 4	
	PG 4	
	0	packets
	0	bytes
	0	queue depth
	0	no buffer discard
	0	shared buffer discard
Related Commands		
Note	This command is only supported on Spectrum™ based switch systems	

show interfaces ethernet counters pfc prio

show interfaces ethernet <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	
Role	admin	
Example	<pre>switch (config) # show interfaces ethernet 1/1 counters pfc prio 1 PFC 1 Rx 0 pause packets 0 pause duration Tx 0 pause packets 0 pause duration</pre>	
Related Commands		
Note	This command is only supported on Spectrum™ based switch systems	

5.14.6.2 QoS Rewrite

qos rewrite pc

qos rewrite pc-enable
qos rewrite pc-disable

Enables or 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
Role	admin
Example	switch (config interface ethernet 1/1) # qos rewrite pc-enable
Related Commands	
Notes	This command is only supported on Spectrum™ based switch systems

qos rewrite dscp

qos rewrite dscp-enable
qos rewrite dscp-disable

Enables or 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
Role	admin
Example	switch (config interface ethernet 1/1) # qos rewrite dscp-enable
Related Commands	
Notes	This command is only supported on Spectrum™ based switch systems

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:	Switch priority → PCP,DEI: 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	
Role	admin	
Example	switch (config interface ethernet 1/1) # qos rewrite map switch -priority 11 pcp 7 dei 0	
Related Commands	N/A	
Notes	This command is only supported on Spectrum™ based switch systems	

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
Role	admin
Example	switch (config interface ethernet 1/1) # qos rewrite map switch -priority 5 dscp 40
Related Commands	N/A
Notes	This command is only supported on Spectrum™ based switch systems

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	Default is “preserve” when router is disabled Default is “enable” when router is enabled (Router can be enabled/disabled using the “ip routing” command)	
Configuration Mode	Config	
History	3.6.1002	
Role	admin	
Example	switch (config) # qos ip rewrite pcp enable	
Related Commands	N/A	
Notes	The parameter “preserve” is only supported on Spectrum based switch systems	

5.14.6.3 Queuing and Scheduling (ETS)

dcb ets enable

dcb ets enable
no dcb ets enable

Sets the switch egress scheduling mode to be weighted round robin.
 The no form of the command sets the switch egress scheduling mode to be strict priority.

Syntax Description	N/A										
Default	ETS is enabled										
Configuration Mode	Config										
History	3.1.0000 3.6.1002 Updated Note										
Role	admin										
Example	<pre>switch (config)# dcb ets enable switch (config)# show dcb ets</pre> ETS enabled <table> <thead> <tr> <th>TC</th><th>Bandwidth</th></tr> </thead> <tbody> <tr><td>0</td><td>25%</td></tr> <tr><td>1</td><td>25%</td></tr> <tr><td>2</td><td>25%</td></tr> <tr><td>3</td><td>25%</td></tr> </tbody> </table> Number of Traffic Class: 4 <pre>switch (config) #</pre>	TC	Bandwidth	0	25%	1	25%	2	25%	3	25%
TC	Bandwidth										
0	25%										
1	25%										
2	25%										
3	25%										
Related Commands	show dcb ets										
Note	The show command output is from a SwitchX® based switch systems										

dcb ets tc bandwidth

dcb ets tc bandwidth <tc-0> <tc-1> <tc-2> <tc-3>
no dcb ets tc bandwidth

Configures the bandwidth limit of the traffic class.
The no form of the command sets the bandwidths per traffic class back to its default.

Syntax Description	tc-i	0-100.										
Default	25% per traffic class											
Configuration Mode	Config											
History	3.1.0000 3.6.1002 Updated Note											
Role	admin											
Example	<pre>switch (config)# dcb ets tc bandwidth 20 20 30 30 switch (config) # show dcb ets</pre> ETS enabled <table><thead><tr><th>TC</th><th>Bandwidth</th></tr></thead><tbody><tr><td>0</td><td>20%</td></tr><tr><td>1</td><td>20%</td></tr><tr><td>2</td><td>30%</td></tr><tr><td>3</td><td>30%</td></tr></tbody></table> Number of Traffic Class: 4 <pre>switch (config) #</pre>		TC	Bandwidth	0	20%	1	20%	2	30%	3	30%
TC	Bandwidth											
0	20%											
1	20%											
2	30%											
3	30%											
Related Commands	show dcb ets											
Note	- The sum of all traffic class bandwidth must be equal to 100 - This command is only supported on SwitchX® based switch systems											

vlan map-priority

vlan map priority <priority> traffic-class <tc>
no vlan map priority <priority>

Maps an VLAN user priority to a traffic class.
 The no form of the command sets the mapping back to default.

Syntax Description	N/A
Default	Priority 0,1 mapped to tc 0 Priority 2,3 mapped to tc 1 Priority 4,5 mapped to tc 2 Priority 6,7 mapped to tc 3
Configuration Mode	Config Interface Ethernet
History	3.1.0000 3.6.1002 Updated Note
Role	admin
Example	switch (config interface ethernet 1/1) # vlan map-priority 1 traffic-class 2
Related Commands	show dcb ets interface
Note	This command is only supported on SwitchX® based switch systems

show dcb ets (SwitchX)

show dcb ets

Displays ETS configuration and operational data.

Syntax Description	N/A										
Default	ETS is enabled.										
Configuration Mode	Any Command Mode										
History	3.1.0000 3.6.1002 Updated Note										
Role	admin										
Example	<pre>switch (config)# show dcb ets</pre> ETS enabled <table> <thead> <tr> <th>TC</th><th>Bandwidth</th></tr> </thead> <tbody> <tr><td>0</td><td>25%</td></tr> <tr><td>1</td><td>25%</td></tr> <tr><td>2</td><td>25%</td></tr> <tr><td>3</td><td>25%</td></tr> </tbody> </table> Number of Traffic Class: 4	TC	Bandwidth	0	25%	1	25%	2	25%	3	25%
TC	Bandwidth										
0	25%										
1	25%										
2	25%										
3	25%										
Related Commands											
Note	The show command output is from a SwitchX® based switch system										

show dcb ets interface

show dcb ets interface <type> <number>

Displays ETS configuration and operational data, per interface.

Syntax Description	type	ethernet or port-channel
	number	interface number, i.e. 1/1
Default	ETS is enabled.	
Configuration Mode	Any Command Mode	
History	3.1.0000	
Role	admin	

Example

```
switch (config)# show dcb ets interface ethernet 1/1
```

```
ETS Port Mode           :ON MODE
ETS Oper State          :INIT STATE
ETS State Machine Type  :Assymetric
-----
```

```
ETS Local Port Info
-----
```

```
TC bandwidth table
-----
```

TC	Bandwidth	RecomBandwidth
0	25%	25%
1	25%	25%
2	25%	25%
3	25%	25%

```
priority assignment table
-----
```

Priority	TC
0	0
1	0
2	1
3	1
4	2
5	2
6	3
7	3

```
Number of Traffic Class: 4
```

```
Willing Status: Disable
-----
```

```
ETS Admin Port Info
-----
```

TC	Bandwidth	RecomBandwidth
0	30%	30%
1	30%	30%
2	30%	30%
3	10%	10%

```
ETS Remote Port Info
-----
```

```
No Remote Entry is Present
-----
```

```
switch (config) #
```

Related Commands

Note

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
Role	admin
Example	switch (config interface ethernet 1/1 traffic-class 0) # bind switch-property 1
Related Commands	N/A
Notes	• Context is egress interface traffic class • This command is only supported on Spectrum™ based switch systems

bandwidth guaranteed

bandwidth guaranteed [<rate>]
no bandwidth guaranteed [<rate>]

Configures the minimum bandwidth for outbound traffic.

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	
Role	admin	
Example	switch (config interface ethernet 1/1 traffic-class 0) # bandwidth guaranteed 0.4G	
Related Commands	N/A	
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 - This command is only supported on Spectrum™ based switch systems	

bandwidth shape

bandwidth shape [<rate>]
no bandwidth shape [<rate>]

Configures the bandwidth shaper for outbound traffic.

Syntax Description	rate Rate in GbE Range: 0 - max speed supported
Default	Maximum port rate (100GbE on Spectrum™ based switches)
Configuration Mode	Config Interface Ethernet Traffic Class Config Interface Port Channel Traffic Class Config Interface MLAG Port Channel Traffic Class
History	3.6.1002
Role	admin
Example	switch (config interface ethernet 1/1 traffic-class 7) # bandwidth shape 0.4G
Related Commands	N/A
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 - This command is only supported on Spectrum™ based switch systems

dcb ets

```
dcb ets [strict | wrr <weight>]
no dcb ets [strict | wrr <weight>]
```

Configures ETS mode to strict or WRR.

Syntax Description	weight
Default	Default is WRR with the following default weights. Traffic class to weight mapping: 0 → 12 1 → 13 2 → 12 3 → 13 4 → 12 5 → 13 6 → 12 7 → 13
Configuration Mode	Config Interface Ethernet Traffic Class Config Interface Port Channel Traffic Class Config Interface MLAG Port Channel Traffic Class
History	3.6.1002
Role	admin
Example	switch (config interface ethernet 1/1 traffic-class 1) # dcb ets wrr 50
Related Commands	N/A
Notes	- Context is egress interface traffic class - This command is only supported on Spectrum™ based switch systems

mc-unaware tc binding

mc-unaware tc binding
no mc-unaware tc binding

Configures the MC-unaware TC binding.
 The no form of the command disables MC-unaware TC binding.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config Interface Ethernet Config Interface Port Channel Config Interface MLAG Port Channel
History	3.6.1002
Role	admin
Example	switch (config interface ethernet 1/1) # mc-unaware tc binding
Related Commands	N/A
Notes	- When the no form is configured, the multicast traffic of a switch-priority that is mapped to TC X is re-mapped to TC X+8 - Context is egress interface - This command is only supported on Spectrum™ based switch systems

show dcb ets (Spectrum)

show dcb ets [interface {ethernet | mlag-port-channel | port-channel} <number>]

Displays ETS information.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Any Command Mode
History	3.6.1002
Role	admin

Example

```
switch (config) # show dcb ets interface ethernet 1/1
Eth1/1
Interface Bandwidth Shape [Mbps]: 100000
Multicast unaware mapping : disabled
```

```
ETS per TC :
TC Scheduling Mode Weight Weight (%)
-----
0 WRR                12      12
1 WRR                13      13
2 WRR                12      12
3 WRR                13      13
4 WRR                12      12
5 WRR                13      13
6 WRR                12      12
7 WRR                13      13
```

```
Bandwidth Shape per TC:
TC Bandwidth Shape [Mbps]
-----
0 100000
1 100000
2 100000
3 100000
4 100000
5 100000
6 100000
7 100000
```

```
Bandwidth Guarantee per TC:
TC Bandwidth Guaranteed [Mbps]
-----
0 0
1 0
2 0
3 0
4 0
5 0
6 0
7 0
```

```
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 N/A

Notes The show command output is from a Spectrum™ based switch systems

5.14.6.4 RED & ECN

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 minimum-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	
Role	admin	
Example	switch (config interfaces ethernet 1/1)# traffic-class 0 congestion-control both minimum-relative 50 maximum-relative 80	
Related Commands		
Note		

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
Role	admin
Example	<pre> switch (config)# show interfaces 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 switch (config) # </pre>

Related Commands

Note

5.15 Access Control List

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, e.g L2/L3 source and destination addresses, protocol, VLAN ID and priority or TCP port.

5.15.1 Configuring Access Control List

Access Control List (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:**

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

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

Step 3. Create a MAC / IPv4 ACL (access-list) entity.

```
switch (config) mac access-list mac-acl
switch (config mac access-list mac-acl) #
```

Step 4. Add a MAC / IP rules to the appropriate access-list.

```
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
switch (config mac access-list mac-acl) #
```

Step 5. Bind the created access-list to an interface (slot/port or port-channel).

```
switch (config)
switch (config) # interface ethernet 1/1
switch (config interface ethernet 1/1) # mac port access-group mac-acl
```

5.15.2 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:**

Step 1. Create access-list action profile:

Step 1a. Create an action access-list profile using the command `access-list action <action-profile-name>`.

Step 1b. Add rule to map a VLAN using the command `vlan-map <vlan-id>` within the action profile configuration mode.

Step 1c. Add action on a rule to strip the VLAN from a packet using the command `vlan-pop` within the action profile configuration mode.

- Step 1d.** Add action on a rule to append a VLAN to a packet using the command `vlan-push` within the action profile configuration mode.
- Step 2.** Create an access-list and bind the action rule:
- a. Create an access-list profile using the command `ipv4/mac access-list`
 - b. Add access list rule using the command `deny/permit (action <action profile name>)`
- Step 3.** Bind the access-list to an interface using the command `ipv4/mac 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
```


5.15.3 Commands

ipv4/mac access-list

```
{ipv4 | mac} access-list <acl-name>
no {ipv4 | mac} access-list <acl-name>
```

Creates a MAC or IPv4 ACL and enter the ACL configuration mode.
The no form of the command deletes the ACL.

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	
Role	admin	
Example	switch (config)# mac access-list my-mac-list switch (config mac access-list my-mac-list)#	
Related Commands	ipv4/port access-group	
Note		

ipv4/mac port access-group

```
{ipv4 | mac} port access-list <acl-name>
no {ipv4 | mac} port access-list <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	
History	3.1.1400	
	3.3.4500	Added MLAG port-channel configuration mode
Role	admin	
Example	switch (config interface ethernet 1/1) # mac port access-group my-list switch (config interface ethernet 1/1) #	
Related Commands	ipv4/mac access-list	
Note	The access control list should be defined prior to the binding action.	

deny/permit (MAC ACL rule)

[seq-number <sequence-number>] {deny|permit} {any | <source-mac> [mask <mac>]} {any | <destination-mac> [mask <mac>]} [protocol <protocol>] [cos <cos-value>] [vlan <vlan-id> | vlan-mask <vlan-mask>] [action <action-id>] 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. The range is:1-65535.
	deny permit	Determines the type of the rule, denies or permits action.
	{any <source-mac> [mask <mac>]}	Sets source MAC and optionally sets a mask for that MAC. The “any” option will cause the rule not to check the source MAC.
	{any <destination-mac> [mask <mac>]}	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. Possible range is: 0x0000-0xffff.
	cos-value	Sets the COS (priority bits) field, possible range is: 0-7.
	vlan-id	Sets the VLAN ID field, possible range is 0-4095
	vlan-mask <vlan-mask>	Sets VLAN group. Range: 0x0000-0x0FFF.
	action	Action name (free string)
Default	No rule is added by default to access control list. Default sequence number is in multiple 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
Role	admin	
Example	<pre>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 switch (config mac access-list my-list) #</pre>	

Related Commands	ipv4/mac access-list ipv4/mac port access-group

Note

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>]
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. The range is:1-65535.
	deny permit	Determines the type of the rule, deny or permit action.
	{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. Valid mask values fall in the range 0-255.
	{any <destination-ip> [mask <ip>]}	Sets destination IP and optionally sets a mask for that MAC. The “any” option causes the rule to not check the destination MAC.
Default	No rule is added by default to access control list. Default sequence number is in multiple of 10.	
Configuration Mode	Config IPv4 ACL	
History	3.1.1400	First version
	3.3.4302	Updated syntax description of mask <ip> parameter
	3.5.1000	Updated seq-number parameter
Role	admin	
Example	switch (config ipv4 access-list my-list) # seq-number 51 deny ip 1.1.1.1 mask 123.12.13.53 45.45.45.0 mask 123.132.21.123 switch (config ipv4 access-list my-list) #	
Related Commands	ipv4/mac access-list ipv4/mac port access-group	
Note		

deny/permit (IPv4 TCP/UDP/ICMP ACL rule)

```
[seq-number <sequence-number>] {permit | deny} {tcp | udp | icmp} {<source-
ip> [mask <ip>] | [any]} [<dest-ip> [mask <ip>] | [any]] [eq-source <port-num-
ber>] [eq-destination <port-number>] [action <action-id>] [eq-code <icmp-
code>] [eq-type <icmp-type>]
no <sequence-number>
```

Creates a rule for IPv4 UDP/TCP/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. The range is:1-65535.
	deny permit	Determines the type of the rule, deny or permit action.
	tcp udp icmp	UDP, TCP, or ICMP rule transport type.
	{any <source-ip> [mask <ip>]}	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.
	{any <destination-ip> [mask <ip>]}	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-source <port-number>]	TCP/UDP/ICMP source port number Range is 0-65535
	[eq-destination <port-number>]	TCP/UDP/ICMP destination port number Range is 0-65535
	eq-code <icmp-code>	Range: 0-255
	eq-type <icmp-type>	Range: 0-255
Default	No rule is added by default to access control list Default sequence number is in multiple of 10	
Configuration Mode	Config IPv4 ACL	
History	3.1.1400	
	3.5.1000	Updated seq-number parameter
	3.6.2002	Added ICMP options and Notes section
Role	admin	
Example	switch (config ipv4 access-list my-list) # seq-number 10 deny tcp any any eq-source 1200 switch (config ipv4 access-list my-list) #	

Related Commands

ipv4/mac access-list
ipv4/mac 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.
 - The parameters “eq-source” and “eq-destination” are not applicable with ICMP
-
-

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
Role	admin
Example	<pre>switch (config)# access-list action my-action switch (config access-list action my-action)# show access-list action my-action Access-list Action my-action Mapped_Vlan_ID Mapped_port Counter_set Policer_ID ===== N/A N/A N/A N/A switch (config access-list action my-action)#</pre>
Related Commands	
Note	

vlan-map

vlan-map <vlan-id>
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	vlan-id	0-4095.
Default	N/A	
Configuration Mode	Config ACL Action	
History	3.2.0230	
Role	admin	
Example	<pre>switch (config access-list action my-action)# vlan-map 10 switch (config access-list action my-action)# show access-list action my-action Access-list Action my-action Mapped_Vlan_ID Mapped_port Counter_set Policer_ID ===== 10 N/A N/A N/A switch (config access-list action my-action)#</pre>	
Related Commands		
Note		

vlan-pop

vlan-pop

Pops VLAN frames from traffic.

Syntax Description	vlan-id	VLAN ID: 0-4095.
Default	N/A	
Configuration Mode	Config ACL Action	
History	3.4.3000	
Role	admin	
Example	<pre>switch (config access-list action my-action)# vlan-pop switch (config access-list action my-action)# show access-list action my-action Access-list Action my-action Popped_Vlan_ID Mapped_port Counter_set Policer_ID ===== N/A N/A N/A N/A switch (config access-list action my-action) #</pre>	
Related Commands		
Note		

vlan-push

vlan-push <vlan-id>

Pushes (or adds) VLAN frames to traffic.

Syntax Description	vlan-id	VLAN ID: 0-4095
Default	N/A	
Configuration Mode	Config ACL Action	
History	3.4.3000	
Role	admin	
Example	<pre>switch (config access-list action my-action)# vlan-push 10 switch (config access-list action my-action)# show access-list action my-action Access-list Action my-action Mapped_Vlan_ID Mapped_port Counter_set Policer_ID ===== 10 N/A N/A N/A switch (config access-list action my-action)#</pre>	
Related Commands		
Note		

show access-list action

show access-list action {<action-profile-name> | **summary**}

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	
Role	admin	
Example	<pre> witch (config)# show access-list action my-action Access-list Action my-action Mapped_Vlan_ID Mapped_port Counter_set Policer_ID ===== 10 N/A N/A N/A switch (config)# </pre>	
Related Commands		
Note		

show mac/ipv4 access-lists

show [mac | ipv4 |] access-lists <access-list-name>

Displays the list of rules for the MAC/IPv4 ACL.

Syntax Description	ipv4 mac	IPv4 or MAC - access list.
	access-list-name	ACL name.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.1.1400	
History	3.3.4500	Updated output
Role	admin	
Example	<pre>switch (config mac access-list my-list) # show mac access-lists my-list mac access-list my-list seq-number p/d smac dmac protocol cos vlan vlan-mask action ===== 10 deny any any 0800 3 3 0x0FFF none 20 deny any any 80 2 6 0x0000 none 30 deny any any any any any 0x0ACB none 40 deny any any any any any N/A none switch (config mac access-list my-list) #</pre>	
Related Commands	deny/permit (MAC ACL rule) deny/permit (IPv4 ACL rule) deny/permit (IPv4 TCP/UDP ACL rule) ipv4/mac access-list ipv4/mac port access-group	
Note		

show mac/ipv4 access-lists summary

show [mac |ipv4 |] access-lists summary

Displays the summary of number of rules per ACL, and the interfaces attached.


Syntax Description	ipv4 mac	IPv4 or MAC - Access list
	access-list-name	ACL name
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.1.1400	
Role	admin	
Example	<pre>switch (config) # show mac access-lists summary mac access-list my-list Total ACEs Configured: 2 Configured on interfaces: Ethernet 1/1 Ethernet 1/2 switch (config) #</pre>	
Related Commands	<pre>deny/permit (MAC ACL rule) deny/permit (IPv4 ACL rule) deny/permit (IPv4 TCP/UDP ACL rule) ipv4/mac access-list ipv4/mac port access-group</pre>	
Note		

5.16 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.

Figure 29 provides an overview of the mirroring functionality.

Figure 29: Overview of 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.

5.16.1 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.

Figure 30: Mirror to Analyzer Mapping



A mirroring session is a monitoring configuration mode that has the following parameters:

Table 61 - Mirroring 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

5.16.1.1 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.

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.

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.

5.16.1.2 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. SwitchX® platforms support up to 7 analyzer ports, and Spectrum™ platforms support up to 2 only, 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.

LAG

The destination interface cannot be a member of LAG when the header format is local.

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.

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).

5.16.1.3 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 in 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.

Figure 31: Header Format Options



5.16.1.4 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, SwitchX 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 SwitchX 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.

5.16.1.5 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.

5.16.2 Configuring Mirroring Sessions

Figure 32 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 con-

nectivity 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.

Figure 32: Mirroring Session



➤ **To configure a mirroring session:**

Step 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.

Step 2. Add source interface(s). Run:

```
switch (config monitor session 1) # add source interface ethernet 1/1 direction both
```

Step 3. Add destination interface. Run:

```
switch (config monitor session 1) # destination interface ethernet 1/2
```

Step 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 traffic-class 2
```



For remote connectivity use the header formats add-vlan or add-ethernet-header. For local connectivity, use local.

Step 5. (Optional) Truncate the mirrored traffic to 64-byte packets. Run:

```
switch (config monitor session 1) # truncate
```

Step 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` option uses flow control to regulate the traffic from the source interfaces.



If the option `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.

Step 7. Enable the session. Run:

```
switch (config monitor session 1) # no shutdown
```

5.16.3 Verifying Mirroring Sessions

➤ **To verify the attributes of a specific mirroring session:**

```
switch (config) # show monitor session 1
Admin: Enable
Status: Up
Truncate: Enable
Destination interface: eth1/2
Congestion type: pause-excessive-frames
Header format: add-ethernet-header
- traffic class 2
- vlan 10
- priority 5
- destination-mac 00:0d:ec:f1:a9:c8

Source interfaces
Interface direction
-----
eth1/1      both
```

➤ **To verify the attributes of running mirroring sessions:**

```
switch (config) # show monitor session summary
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
```

5.16.4 Commands

5.16.4.1 Config

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 is: • 1-7 for SwitchX® • 1-2 for Spectrum™
Default	N/A	
Configuration Mode	Config	
History	3.3.3500	
Role	admin	
Example	switch (config)# monitor session 1 switch (config) monitor session 1)#	
Related Commands		
Note		

5.16.4.2 Config Monitor Session

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 <type> <number>	Sets the interface type and number (e.g. ethernet 1/2)
	force	The user does not need to shutdown the port prior the operation.
Default	no destination interface	
Configuration Mode	Config Monitor Session	
History	3.3.3500	First version
	3.3.4100	Added force argument
Role	admin	
Example	switch (config monitor session 1) # destination interface ethernet 1/2 switch (config monitor session 1)#	
Related Commands		
Note		

shutdown

shutdown
no shutdown

Disables the session.
The no form of the command enables the session.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config Monitor Session
History	3.3.3500
Role	admin
Example	switch (config monitor session 1) # no shutdown switch (config monitor session 1)#
Related Commands	
Note	

add source interface

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 <type> <number>	Configures interface as “ethernet” or “port-channel”.
	direction <d-type>	Configures the direction of the mirrored traffic. The options are as follows: • egress – sets the egress traffic to be monitored • ingress – sets the ingress traffic to be monitored • both – sets egress and ingress traffic to be monitored
Default	N/A	
Configuration Mode	Config Monitor Session	
History	3.3.3500	
	3.5.1000	Updated
Role	admin	
Example	switch (config monitor session 1) # add source interface ethernet 1/1 direction ingress switch (config monitor session 1)#	
Related Commands		
Note	• 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” • Spectrum™ based switch systems only support mirroring ingress traffic	

header-format

header-format {local [traffic-class <tc>] | add-vlan <vlan-id> [priority <prio>] [traffic-class <tc>] [switch-priority <sp>] | add-ethernet-header destination-mac <mac-address> [add-vlan <vlan-id> [priority <prio>]] [traffic-class <tc>]}
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.
	traffic-class <tc>	Changes the egress traffic class of the frame. Range: 0-3.
	switch-priority <sp>	Changes the egress switch priority of the frame. Range: 0-15.
	add-vlan <vlan-id>	An 802.1q VLAN tag is added to the frame.
	priority <prio>	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
Role	admin	
Example	switch (config monitor session 1) # header-format add-ethernet-header destination-mac 00:0d:ec:f1:a9:c8 add-vlan 10 priority 5 traffic-class 2 switch (config monitor session 1)#	

Related Commands

Note

- If add-ethernet-header is used, the source MAC address is the one attached to the switch
 - The parameter traffic-class is only available on SwitchX® based switch systems
 - The parameter switch-priority is only available on Spectrum™ based switch systems
-

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
Role	admin
Example	switch (config monitor session 1) # truncate switch (config monitor session 1)#
Related Commands	
Note	This command applies for all sessions on the same analyzer port.

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	
	3.3.4000	Added Syntax Description.
Role	admin	
Example	<pre>switch (config monitor session 1) # congestion pause-excessive-frames switch (config monitor session 1)#</pre>	
Related Commands		
Note	This command applies for all sessions on the same analyzer port.	

5.16.4.3 Show

show monitor session

show monitor session <session-id>

Displays monitor session configuration and status.

Syntax Description	session-id	The monitor session ID. Range is 1-7.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.3500	
	3.5.1000	Updated Note section
Role	admin	
Example	<pre>switch (config) # show monitor session 1 Admin: Enable Status: Up Truncate: Enable Destination interface: eth1/2 Congestion type: pause-excessive-frames Header format: add-ethernet-header - traffic class 2 - vlan 10 - priority 5 - destination-mac 00:0d:ec:f1:a9:c8 Source interfaces Interface direction ----- eth1/1 both switch (config) #</pre>	
Related Commands		
Note	The output provided is from a SwitchX® based switch system.	

show monitor session summary

show monitor session summary

Displays monitor session configuration and status summary.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.3500
Role	admin
Example	<pre>switch (config) # show monitor session summary 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 switch (config) #</pre>
Related Commands	
Note	


5.17 sFlow

sFlow (ver. 5) is a procedure for statistical monitoring of traffic in networks. MLNX-OS 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.

Figure 33: sFlow Functionality Overview



5.17.1 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.

5.17.2 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 following statistics are gathered by the CPU:

Table 62 - List of Statistical Counters

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.

5.17.3 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.

5.17.4 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.

5.17.5 Configuring sFlow

➤ **To configure the sFlow agent:**

Step 1. Unlock the sFlow commands. Run:

```
switch (config) # protocol sflow
```

Step 2. Enable sFlow on the system. Run:

```
switch (config) # sflow enable
```

Step 3. Enter sFlow configuration mode. Run:

```
switch (config) # sflow
switch (config sflow) #
```


Step 4. Set the central collector's IP. Run:

```
switch (config sflow) # collector-ip 10.10.10.10
```

Step 5. Set the agent-ip used in the sFlow header. Run:

```
switch (config sflow) # agent-ip 20.20.20.20
```

Step 6. (Optional) Set the sampling rate of the mechanism. Run:

```
switch (config sflow) # sampling-rate 16000
```



This means that one every 16000 packet gets collected for sampling.

Step 7. (Optional) Set the maximum size of the data path sample. Run:

```
switch (config sflow) # max-sample-size 156
```

Step 8. (Optional) Set the frequency in which counters are polled. Run:

```
switch (config sflow) # counter-poll-interval 19
```

Step 9. (Optional) Set the maximum size of the datagrams sent to the central collector. Run:

```
switch (config sflow) # max-datagram-size 1500
```

Step 10. Enable the sFlow agent on the desired interfaces. Run:

```
switch (config interface ethernet 1/1)# sflow enable  
switch (config interface port-channel 1)# sflow enable
```

5.17.6 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
```

5.17.7 Commands

5.17.7.1 Config

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
Role	admin
Example	switch (config) # protocol sflow switch (config) #
Related Commands	
Note	

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
Role	admin
Example	switch (config) # sflow enable switch (config) #
Related Commands	
Note	

sflow

sflow

Enters sFlow configuration mode.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.3.3500
Role	admin
Example	switch (config) # sflow switch (config sflow) #
Related Commands	
Note	

5.17.7.2 Config sFlow

sampling-rate

sampling-rate <rate>
no sampling-rate

Sets sFlow sampling ratio.
The no form of the command resets this parameter to its default value.

Syntax Description	rate	Sets the number of packets passed before selecting one for sampling. The range is 4000-16777215. Zero disables sampling.
Default	16000	
Configuration Mode	Config sFlow	
History	3.3.3500	
Role	admin	
Example	switch (config sflow) # sampling-rate 16111 switch (config sflow) #	
Related Commands		
Note		

max-sample-size

max-sample-size <packet-size>
no max-sample-size

Sets 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. The range is 64-256 bytes.
Default	128 bytes	
Configuration Mode	Config sFlow	
History	3.3.3500	
Role	admin	
Example	switch (config sflow) # max-sample-size 165 switch (config sflow) #	
Related Commands		
Note	Sampled payload beyond the configured size is discarded.	

counter-poll-interval

counter-poll-interval <seconds>
no counter-poll-interval

Sets 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 is 5-3600 seconds. Zero disables the statistic polling.
Default	20 seconds	
Configuration Mode	Config sFlow	
History	3.3.3500	
Role	admin	
Example	switch (config sflow) # counter-poll-interval 30 switch (config sflow) #	
Related Commands		
Note		

max-datagram-size

max-datagram-size <packet-size>
no max-datagram-size

Sets 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. The range is 200-9216 bytes.
Default	1400 bytes	
Configuration Mode	Config sFlow	
History	3.3.3500	
Role	admin	
Example	<pre>switch (config sflow) # max-datagram-size 9216 switch (config sflow) #</pre>	
Related Commands		
Note	This packet contains the data sample as well as the statistical counter data.	

collector-ip

**collector-ip <ip-address> [udp-port <udp-port-number>]
no collector-ip [<ip-address> udp-port]**

Sets 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 <udp-port-number>	Sets the collector UDP port number.
Default	ip-address: 0.0.0.0 udp-port-number: 6343	
Configuration Mode	Config sFlow	
History	3.3.3500	
Role	admin	
Example	switch (config sflow) # collector-ip 10.10.10.10 switch (config sflow) #	
Related Commands		
Note		

agent-ip

agent-ip {<ip-address> | interface [ethernet <slot/port> | port-channel <channel-group>] | <if-name> | loopback <number> | vlan <id>}
no agent-ip

Sets 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/port-channel 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 <number>	Loopback interface number. Range: 1-32.
	vlan <id>	Interface VLAN. Range: 1-4094.
Default	ip-address: 0.0.0.0	
Configuration Mode	Config sFlow	
History	3.3.3500	
	3.3.5200	Updated "interface" parameters
Role	admin	
Example	switch (config sflow) # agent-ip 20.20.20.20 switch (config sflow) #	
Related Commands		
Note	The IP address here is used in the sFlow header.	

clear counters

clear counters

Clears sFlow counters.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config sFlow
History	3.3.3500
Role	admin
Example	switch (config sflow) # clear counters switch (config sflow) #
Related Commands	
Note	

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 MLAG port-channel configuration mode
Role	admin
Example	switch(config interface ethernet 1/1)# sflow enable ... switch(config interface port-channel 1)# sflow enable
Related Commands	
Note	

5.17.7.3 Show

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 output example
Role	admin
Example	<pre>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</pre>
Related Commands	
Note	

5.18 Transport Applications

5.18.1 RDMA over Converged Ethernet (RoCE)


5.18.1.1 RoCE Overview

Remote Direct Memory Access (RDMA) is the remote memory management capability that allows server to server data movement directly between application memory without any CPU involvement. RDMA over Converged Ethernet (RoCE) is a mechanism to provide this efficient data transfer with very low latencies on loss-less Ethernet networks. With advances in data center convergence over reliable Ethernet, ConnectX® EN with RoCE uses the proven and efficient RDMA transport to provide the platform for deploying RDMA technology in mainstream data center application at 10GigE and 40GigE link-speed. ConnectX® EN with its hardware offload support takes advantage of this efficient RDMA transport services over Ethernet to deliver ultra-low latency for performance-critical and transaction intensive applications such as financial, database, storage, and content delivery networks. RoCE encapsulates IB transport and GRH headers in Ethernet packets bearing a dedicated ether type. While the use of GRH is optional within subnets, it is mandatory when using RoCE. Applications written over IB verbs should work seamlessly, but they require provisioning of GRH information when creating address vectors. The library and driver are modified to provide mapping from GID to MAC addresses required by the hardware.

5.18.1.1.1 IP Routable (RoCEv2)

A straightforward extension of the RoCE protocol enables traffic to operate in layer 3 environments. This capability is obtained via a simple modification of the RoCE packet format. Instead of the GRH used in RoCE, routable RoCE packets carry an IP header which allows traversal of IP L3 Routers and a UDP header that serves as a stateless encapsulation layer for the RDMA Transport Protocol Packets over IP.

Figure 34: RoCEv2 and RoCE Frame Format Differences




The proposed RoCEv2 packets use a well-known UDP destination port value that unequivocally distinguishes the datagram. Similar to other protocols that use UDP encapsulation, the UDP source port field is used to carry an opaque flow-identifier that allows network devices to imple-

ment packet forwarding optimizations (e.g. ECMP) while staying agnostic to the specifics of the protocol header format.

Furthermore, since this change exclusively affects the packet format on the wire, and due to the fact that with RDMA semantics packets are generated and consumed below the AP applications can seamlessly operate over any form of RDMA service (including the routable version of RoCE as shown in Figure 2), in a completely transparent way¹.

Figure 35: RoCEv2 Protocol Stack



5.18.1.2 RoCE Configuration

In order to function reliably, RoCE requires a form of flow control. While it is possible to use global flow control, this is normally undesirable, for performance reasons.

The normal and optimal way to use RoCE is to use Priority Flow Control (PFC). To use PFC, it must be enabled on all endpoints and switches in the flow path.

In the following section we present instructions to configure PFC on Mellanox ConnectX™ cards. There are multiple configuration steps required, all of which may be performed via PowerShell. Therefore, although we present each step individually, you may ultimately choose to write a PowerShell script to do them all in one step. Note that administrator privileges are required for these steps.

For further information, please refer to the following URL:

<http://blogs.technet.com/b/josebda/archive/2012/07/31/deploying-windows-server-2012-with-smb-direct-smb-over-rdma-and-the-mellanox-connectx-3-using-10gbe-40gbe-roce-step-by-step.aspx>

1. Standard RDMA APIs are IP based already for all existing RDMA technologies

5.18.1.2.1 Prerequisites

The following are the driver's prerequisites in order to set or configure RoCE:

- ConnectX®-3 and ConnectX®-3 Pro firmware version 2.30.3000 or higher
- Set HCA to use Ethernet protocol:
Display the Device Manager and expand "System Devices".

5.18.1.2.2 Configuring Windows Host



Since PFC is responsible for flow controlling at the granularity of traffic priority, it is necessary to assign different priorities to different types of network traffic.

As per RoCE configuration, all ND/NDK traffic is assigned to one or more chosen priorities, where PFC is enabled on those priorities.

Configuring Windows host requires configuring QoS.

5.18.1.2.2.1 Using Global Pause Flow Control (GFC)

➤ **To use Global Pause Flow Control (GFC) mode, disable QoS and Priority:**

```
PS $ Disable-NetQosFlowControl
PS $ Disable-NetAdapterQos
```

5.18.1.3 Configuring Switch Systems

➤ **To enable RoCE, the SwitchX should be configured as follows:**

- Ports facing the host should be configured as access ports, and either use global pause or Port Control Protocol (PCP) for priority flow control
- Ports facing the network should be configured as trunk ports, and use Port Control Protocol (PCP) for priority flow control

5.18.1.4 Configuring Router (PFC only)

The router uses L3's DSCP value to mark the egress traffic of L2 PCP. The required mapping, maps the three most significant bits of the DSCP into the PCP. This is the default behavior, and no additional configuration is required.

5.18.1.4.1 Copying Port Control Protocol (PCP) Between Subnets

The captured PCP option from the Ethernet header of the incoming packet can be used to set the PCP bits on the outgoing Ethernet header.

5.18.1.5 Configuring the RoCE Mode

Configuring the RoCE mode requires the following:

- RoCE mode is configured per-driver and is enforced on all the devices in the system



The supported RoCE modes depend on the firmware installed. If the firmware does not support the needed mode, the fallback mode would be the maximum supported RoCE mode of the installed NIC.

RoCE mode can be enabled and disabled via PowerShell.

➤ **To enable RoCE using the PowerShell:**

- Open the PowerShell and run:

```
Set-MlnxDriverCoreSetting -RoceMode 1
```

➤ **To enable RoCEv2 using the PowerShell:**

- Open the PowerShell and run:

```
Set-MlnxDriverCoreSetting -RoceMode 2
```

➤ **To disable any version of RoCE using the PowerShell:**

Open the PowerShell and run:

```
Set-MlnxDriverCoreSetting -RoceMode 0
```

➤ **To check current version of RoCE using the PowerShell:**

Step 1. Open the PowerShell and run:

```
Get-MlnxDriverCoreSetting
```

Step 2. Example output:

```
Caption          : DriverCoreSettingData 'mlx4_bus'
Description      : Mellanox Driver Option Settings
.
.
.
RoceMode         : 0
```

5.19 802.1x Protocol

The 802.1x 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. Mellanox 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.

5.19.1 802.1x 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 unau-

thenticated. 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.

5.19.2 Configuring 802.1x

➤ To configure 802.1x on the switch

Step 1. Enable 802.1x protocol. Run:

```
switch (config) # protocol dot1x
```

Step 2. Enable the system as authenticator. Run:

```
switch (config) # dot1x system-auth-control
```

Step 3. Configure RADIUS server parameters. Run:

```
switch (config) # dot1x radius-server host 10.10.10.10 key my4uth3nt1c4t10nk3y retrans-  
mit 2 timeout 3
```

Step 4. Enter the configuration mode of an Ethernet interface. Run:

```
switch (config) # interface ethernet 1/1  
switch (config interface ethernet 1/1) #
```

Step 5. Configure the interface as a port access entity authenticator. Run:

```
switch (config interface ethernet 1/1) # dot1x pae authenticator
```

Step 6. Configure the interface to perform authentication on ingress traffic. Run:

```
switch (config interface ethernet 1/1) # dot1x port-control auto
```

Step 7. Verify 802.1x configuration. Run:

```
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  
switch (config interface ethernet 1/1) #
```

5.19.3 Commands

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
Role	admin
Example	switch (config)# protocol dot1x
Related Commands	
Note	

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
Role	admin
Example	switch (config)# dot1x clear-statistics
Related Commands	
Note	

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
Role	admin
Example	switch (config interface ethernet 1/2)# dot1x system-auth-control
Related Commands	
Note	

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
Role	admin	
Example	switch (config interface ethernet 1/2)# dot1x host-mode single-host	
Related Commands		
Note		

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	
Role	admin	
Example	switch (config interface ethernet 1/2)# dot1x port-control auto	
Related Commands		
Note		

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. The time configured is in seconds. Range: 0-5.
	timeout	Configures 802.1x global RADIUS timeout value for servers. The time configured is in seconds. Range: 1-60.
Default	auth-port: 1812 key: empty string retransmit: 1 timeout: 3	
Configuration Mode	Config	
History	3.4.2008	
Role	admin	
Example	switch (config)# dot1x radius-server host 10.10.10.10 auth-port 65535 prompt-key enable	
Related Commands		
Note	- 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	

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	No re-authentication
Configuration Mode	Config Interface Ethernet
History	3.4.2008
Role	admin
Example	switch (config interface ethernet 1/2)# dot1x reauthenticate
Related Commands	
Note	

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
Role	admin
Example	switch (config)# dot1x system-auth-control
Related Commands	
Note	

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 seconds.
Default	3600 seconds	
Configuration Mode	Config Interface Ethernet	
History	3.4.2008	
Role	admin	
Example	switch (config interface ethernet 1/2)# dot1x timeout reauthentication 3600	
Related Commands		
Note		

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 seconds.
Default	60 seconds	
Configuration Mode	Config Interface Ethernet	
History	3.4.2008	
Role	admin	
Example	switch (config interface ethernet 1/2)# dot1x timeout quiet-period 60	
Related Commands		
Note		

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 seconds.
Default	30 seconds	
Configuration Mode	Config Interface Ethernet	
History	3.4.2008	
Role	admin	
Example	switch (config interface ethernet 1/2)# dot1x timeout quiet-period 30	
Related Commands		
Note		

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	
Role	admin	
Example	switch (config interface ethernet 1/2)# dot1x max-req 2	
Related Commands		
Note		

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
Role	admin
Example	<pre>switch (config)# show dot1x</pre> System authentication is enabled <pre> ----- 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 ... switch (config)# </pre>
Related Commands	
Note	

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	
Role	admin	
Example	<pre>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 switch (config interface ethernet 1/2)#</pre>	
Related Commands		
Note		

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	
Role	admin	
Example	<pre>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 switch (config)#</pre>	
Related Commands		
Note		

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
Role	admin
Example	<pre>switch (config)# show dot1x radius 802.1x RADIUS defaults: Key: ***** Timeout: 3 Retransmit: 1 No 802.1x RADIUS servers configured. switch (config)#</pre>
Related Commands	
Note	

5.20 Priority Flow Control

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:**

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

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

Step 3. Enable PFC globally on the switch. Run:

```
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:**

Step 1. Log in as admin.

Step 2. Enter config mode. Run:

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

Step 3. Enable PFC globally on the switch. Run:

```
switch (config) # dcb priority-flow-control enable
# 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) #
```

Step 4. Choose the desirable priority you want to enable using the command `dcb priority-flow-control priority <pri[0..7]> enable`.

```
switch (config) # dcb priority-flow-control priority 5 enable
```

➤ **To enable PFC per interface:**

Step 1. Log in as admin.

Step 2. Change to config mode. Run:

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

Step 3. Enable PFC globally on the switch. Run:

```
switch (config) # dcb priority-flow-control enable
```

Step 4. Choose the desirable priority you want to enable using the command `dcb priority-flow-control priority <pri[0..7]> enable`

```
switch (config) # dcb priority-flow-control 5 enable
```

Step 5. Change to Interface mode. Run:

```
switch (config) #
switch (config) # interface ethernet 1/1
switch (config interface ethernet 1/1) #
```

Step 6. 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.

5.20.1 Flow Control Threshold Configuration for Spectrum

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.

Figure 36: Xon/Xoff Configuration



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 [Section 5.21, “Shared Buffers,”](#) on [page 1005](#) for more information on flow control.

5.20.2 Commands

dcb priority-flow-control enable

dcb priority-flow-control enable [force]
no dcb priority-flow-control enable [force]

Enables PFC globally on the switch.
The no form of the command globally disables PFC on the switch.

Syntax Description	force	Forces operation
Default	PFC is disabled.	
Configuration Mode	Config	
History	3.1.0000	
	3.3.0000	Updated Example
Role	admin	
Example	<pre>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 Priority Enabled List : Priority Disabled List :0 1 2 3 4 5 6 7 TC Lossless --- - 0 N 1 Y 2 Y 3 N Interface PFC admin PFC oper ----- 1/1 Disabled Disabled 1/2 Disabled Disabled 1/3 Disabled Disabled 1/4 Disabled Disabled ... switch (config) #</pre>	
Related Commands	show dcb priority-flow-control	
Note	This command asks the user to approve traffic loss because some interfaces with DCB mode activated might get shut down.	

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
Role	admin
Example	<pre> switch (config)# dcb priority-flow-control priority 0 enable 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 Disabled Disabled 1/2 Disabled Disabled 1/3 Disabled Disabled 1/4 Disabled Disabled ... switch (config) # </pre>
Related Commands	show dcb priority-flow-control
Note	

dcb priority-flow-control mode on

dcb priority-flow-control mode on [force]
no dcb priority-flow-control mode

Enables PFC per interface.
The no form of the command disables PFC per interface.

Syntax Description	force	Force command implementation.																																	
Default	PFC is disabled for all interfaces.																																		
Configuration Mode	Config Interface Ethernet Config Interface Port Channel Config Interface MLAG Port Channel																																		
History	3.1.0000 3.3.4500Added MLAG port-channel configuration mode																																		
Role	admin																																		
Example	<pre>switch (config interface ethernet 1/1) # dcb priority-flow-control mode on switch (config interface ethernet 1/1) # show dcb priority-flow-control</pre> PFC enabled Priority Enabled List : 0 Priority Disabled List : 1 2 3 4 5 6 7 <table><tr><td>TC</td><td>Lossless</td></tr><tr><td>---</td><td>-----</td></tr><tr><td>0</td><td>N</td></tr><tr><td>1</td><td>Y</td></tr><tr><td>2</td><td>Y</td></tr><tr><td>3</td><td>N</td></tr></table> <table><tr><td>Interface</td><td>PFC admin</td><td>PFC oper</td></tr><tr><td>-----</td><td>-----</td><td>-----</td></tr><tr><td>1/1</td><td>On</td><td>Enabled</td></tr><tr><td>1/2</td><td>Disabled</td><td>Disabled</td></tr><tr><td>1/3</td><td>Disabled</td><td>Disabled</td></tr><tr><td>1/4</td><td>Disabled</td><td>Disabled</td></tr><tr><td>...</td><td></td><td></td></tr></table> <pre>switch (config) #</pre>		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	...		
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	show dcb priority-flow-control																																		
Note																																			

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	
Role	admin	
Example	<pre> switch (config interface ethernet 1/1) # 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 ... switch (config) # </pre>	
Related Commands		
Note		

5.21 Shared Buffers



This section is relevant only for Spectrum™ based switch systems.

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.

5.21.1 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 on which the packet is going to be transmitted
- Priority
 - Switch priority (SP) – internal identifier of the packet priority which is used as a key for several internal switch functions and decisions, specifically 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.
 - 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.

Buffers configuration mechanism is providing a way to allocate buffer space for specific traffic types based on the following classification parameters.

- iPort – traffic that arrived on a specific port
- iPort.PG – traffic that arrived on a specific port and mapped to a specific PG
- ePort – traffic that is going to be transmitted on a specific port
- ePort.TC – traffic that is going to be transmitted on a specific port and mapped to a specific TC

By default, multicast packets (including flooding and broadcast) are counted on the egress side. However, multicast packets consume the physical memory space of a single packet and, hence, using native buffering calculations, the multicast packet may negatively affect buffer utilization.

Counting multicast traffic only once is not possible since, unlike unicast traffic where the TC is used as the region indicator of egress traffic, multicast traffic can be transmitted using different TCs on different ports. Therefore, instead of using TC as an egress region indicator, SP is used. Thus, the egress region for multicast traffic is named MC,SP. Hence the following classification parameters for multicast traffic are used.

- MC – traffic to be transmitted as multicast
- MC.SP – traffic to be transmitted as multicast on a specific SP

5.21.2 Buffering 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 (cannot be configured for {iPort}, {ePort}, or {MC})

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). When PG buffer occupancy reaches the threshold, the specific flow control packet is sent.

If there is a physical buffer space for an arriving packet, it is temporarily stored for processing. After processing its egress port, TC and ingress PG are defined. Then, it can be evaluated for eligibility for being stored in the buffer space until it is forwarded.

Buffer eligibility is defined based on the following conditions:

- There is available quota within at least one of the four reserved allocation regions
 - For lossy traffic: $iPort.PG.usage < iPort.PG.reserved \parallel iPort.usage < iPort.reserved \parallel ePort.TC.usage < ePort.TC.reserved \parallel ePort.usage < ePort.reserved$
 - For lossless traffic: $ePort.TC.usage < ePort.TC.reserved \parallel ePort.usage < ePort.reserved$

Note: Ingress check is not performed since all the ingress reserved space is allocated for headroom.

- If a packet is below the all aforementioned four shared allocation thresholds:
 $iPort.PG.usage < iPort.PG.shared \ \&\& \ iPort.usage < iPort.shared \ \&\& \ ePort.TC.usage < ePort.TC.shared \ \&\& \ ePort.usage < ePort.shared$

If a packet is not eligible for buffering:

- For lossy traffic: Packet is dropped
- For lossless traffic: Packet stays in headroom

The eligible packet is counted in usage for the egress regions (ePort, ePort,tc or MC, MC,SP). A packet in lossy traffic is counted for usage in the ingress regions (iPort, iPort,PG). An eligible packet in lossless traffic is counted for usage in the ingress iPort region also, but if it is not eligible and stayed in the headroom, it is counted in its ingress region (iPort,PG) causing it to reach closer to the Xoff threshold.

5.21.3 Pools

Shared buffer space can be statically divided among multiple pools. Each region (iPort, ePort, MC, iPort.PG, ePort.TC and MC.SP) is mapped to specific pools. The pools are divided to ingress pools (iPools) and egress pools (ePools).

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.
Note: 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. It is configured in percentage out of the bounded pool size.
 - 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 $< \alpha$ * free pool space

The port region is counted against the pool that the PG/TC region of the packet is mapped to.

5.21.4 Default Configurations

5.21.4.1 Default Lossy Configuration

The default, out-of-box configuration provides the following settings:

- Pool allocation for ingress control and data packets
 - Each port has a reserved quota and in addition shared buffers

- A single buffer (PG) per port for data packets
- A single buffer (PG) per port for control packets – cannot be configured by the user
- Pool allocation for egress control and data packets
 - Each port has a (small) reserved quota and in addition shared buffers
 - 8 TC per port for data packets
 - A single buffer per port for control packets – cannot be configured by the user
- Pool allocation for egress CPU traffic
 - Each TC has a reserved quota and in addition shared buffers
- Only iPort.PG and ePort.TC enforcement is used, not iPort and ePort enforcement

All the switch-priorities are mapped to ingress PG 0. Each switch-priority i is mapped into a corresponding traffic class i .

5.21.4.2 Default Lossless Configuration

One can switch from lossy to lossless defaults by disabling/enabling global flow control.

The lossless buffer allocation is identical to the lossy default allocation with different shared buffer dynamic thresholds and with an addition of flow control thresholds.

The default Xon and Xoff thresholds are both set to 17KB. The reserved buffer is set to 90KB. It allows having a 100 meter lossless link working at 100GbE, supporting 9KB MTU packets.

5.21.5 Configuration Example

The following example exhibits how to divide the buffer among traffic priorities. Assuming that over an out-of-box lossy default configuration is set, the user here configures buffering configuration for lossless traffic classified to switch-priority 3.

The changes on the default configuration are summarized in the following:

- Ingress:
 - Default reserved PG buffer is reduced from 90KB to 20KB, freeing 70KB for lossless traffic
- Egress:
 - TC3 shared α is configured to infinite, as recommended for TCs with lossless traffic.

Example:

```
// Setting PFC on priority 3
switch (config) # dcb priority-flow-control enable force
switch (config) # dcb priority-flow-control priority 3 enable
switch (config) # interface ethernet <id> dcb priority-flow-control mode on force
// Reducing default PG size
switch (config)# interface ethernet <id> ingress-buffer iPort.pg0 map pool ipool0 type
lossy reserved 20K shared alpha 8
```

```
// Setting lossless ingress buffer PG3 and lossless egress TC3
switch (config)# interface ethernet <id> ingress-buffer iPort.pg3 map pool iPool0 type
lossless reserved 70K xoff 17K xon 17K shared alpha 2
switch (config)# interface ethernet <id> egress-buffer ePort.tc3 map pool ePool0
reserved 4K shared alpha inf
// Mapping switch priority 3 to lossless ingress PG buffer
(config)# interface ethernet <id> ingress-buffer iport.pg3 bind switch-priority 3
```

If the user wants to allocate a separate pool for the new lossless traffic. The changes needed are as follows:

- Ingress:
 - Default reserved PG buffer is reduced from 90KB to 20KB, freeing up more than 70KB for lossless traffic
 - Default pool is reduced from 7960K to 3000K. The rest is allocated to the new pool.
- Egress:
 - TC3 shared alpha is configured to infinite as recommended for TCs with lossless traffic. Default pool is reduced from 14232KB to 4888K. The rest is allocated to the new pool.

Example:

```
// Setting PFC on priority 3
switch (config)# dcb priority-flow-control enable force
switch (config)# dcb priority-flow-control priority 3 enable
switch (config)# interface ethernet <id> dcb priority-flow-control mode on force
// Reducing default PG size
switch (config)# interface ethernet <id> ingress-buffer iPort.pg0 map pool ipool0 type
lossy reserved 20K shared alpha 8
// Setting separate pool for lossless traffic
// Reducing data pool
switch (config)# pool iPool0 direction ingress size 3000 type dynamic
switch (config)# pool ePool0 direction egress size 4888 type dynamic
// Defining lossless pool #1
switch (config)# pool iPool1 direction ingress size 7768 type dynamic
switch (config)# pool ePool1 direction egress size 7768 type dynamic
// Setting lossless ingress buffer PG3 and lossless egress TC3
// Setting iPool1 for infinite alpha
switch (config)# interface ethernet <id> ingress-buffer iPort.pg3 map pool ipool1 type
lossless reserved 70K xoff 17K xon 17K shared alpha 2
switch (config)# interface ethernet <id> ingress-buffer iPort pool iPool1 reserved 0K
shared alpha inf
switch (config)# interface ethernet <id> egress-buffer ePort.tc3 map pool epool1
reserved 4K shared alpha inf
// Mapping switch priority 3 to lossless ingress PG buffer
switch (config)# interface ethernet <id> ingress-buffer iport.pg3 bind switch-priority
3
```




When the egress traffic class (TC) region buffer size exceeds the TX α (max) threshold, the non-eligible packet is dropped (does not stay in the headroom) regardless whether it belongs to a lossy or lossless ingress buffer. Therefore, the recommendation is to map lossless traffic to separate TCs than lossy traffic and to configure egress α (max) threshold of these TCs to infinity in order to avoid dropping lossless traffic.

5.21.6 Commands

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	
Role	admin	
Example	switch (config interface ethernet 1/1)# ingress-buffer iPort.pg1 switch (config interface ethernet 1/1 ingress-buffer iPort.pg1)#	
Related Commands		
Note	iPort.pg9 is reserved for control traffic and hence cannot be edited	

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	
Role	admin	
Example	switch (config interface ethernet 1/1)# egress-buffer ePort.tc4 switch (config interface ethernet 1/1 egress-buffer ePort.tc4)#	
Related Commands		
Note	ePort.tc16 is reserved for control traffic and hence cannot be edited	

pool reserved

pool <pool-name> reserved <reserved> shared {alpha | max} <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, iPool1, iPool2, iPool3
	reserved	Amount of reserved memory for the buffer in bytes
	shared	The amount of shared memory for this buffer - When working in alpha mode, alpha can have the values 0, 1/128, 1/64 ... 1, 2, 4, ... 64, inf - When working in max mode, the shared size is defined as a percentage from the pool size
Default	According to system default OOB configuration	
Configuration Mode	Config Interface Ethernet Egress Buffer Config Interface Ethernet Ingress Buffer	
History	3.6.1002	
Role	admin	
Example	switch (config interface ethernet 1/1 ingress-buffer iPort)# pool iPool0 reserved 90K shared alpha 1/8	
Related Commands		
Note		

map pool

map pool <pool-name> type <type> reserved <reserved> [xoff <xoff> xon [<xon>]
shared {alpha | max} <shared>

Configures the buffer.

The no form of the command resets the values to their default.

Syntax Description	pool-name	Possible values: iPool0, iPool1, iPool2, iPool3
	reserved	Amount of reserved memory for the buffer in bytes
	xoff	Relevant only on lossless type, Xoff threshold in bytes
	xon	Relevant only on lossless type, Xon threshold in bytes
	shared	The amount of shared memory for this buffer - When working in alpha mode, alpha can have the values 0, 1/128, 1/64 ... 1, 2, 4, ... 64, inf - When working in max mode, the shared size is defined as a percentage from the pool size
Default	According to system default OOB configuration	
Configuration Mode	Config Interface Ethernet Egress Buffer Config Interface Ethernet Ingress Buffer	
History	3.6.1002	
Role	admin	
Example	switch (config interface ethernet 1/1 ingress-buffer iPort.pg0)# map pool iPool0 type lossless reserved 90K xoff 17K xon 17K shared alpha 1/8	
Related Commands		
Note		

bind switch-priority

bind switch-priority <list-of-switch-priorities>

Bind a switch priority (SP) to an ingress buffer.
The no form of the command resets the values to their default.

Syntax Description	list-of-switch-priorities Possible values: 0-7
Default	According to system default OOB configuration
Configuration Mode	Config Interface Ethernet Egress Buffer Config Interface Ethernet Ingress Buffer
History	3.6.1002
Role	admin
Example	switch (config interface ethernet 1/1 ingress-buffer iPort.pg1)# bind switch-priority 0 1
Related Commands	
Note	

description

description <description>

Configures buffer description.

The no form of the command resets the values to their default.

Syntax Description	description	Text string
Default	""	
Configuration Mode	Config Interface Ethernet Egress Buffer Config Interface Ethernet Ingress Buffer	
History	3.6.1002	
Role	admin	
Example	switch (config interface ethernet 1/1 ingress-buffer iPort.pg1)# description example	
Related Commands		
Note		

pool direction

pool <pool-name> direction <direction> size <size> type <type>

Configures pool.

The no form of the command resets the values to their default.

Syntax Description	pool	Possible values: iPool0, iPool1, iPool2, iPool3
	direction	Ingress or egress traffic
	size	Size of pool in bytes
	type	Static or dynamic
Default	N/A	
Configuration Mode	Config Interface Ethernet Egress Buffer Config Interface Ethernet Ingress Buffer	
History	3.6.1002	
Role	admin	
Example	switch (config interface ethernet 1/1 ingress-buffer iPort.pg1)# pool iPool1 direction ingress size 1M type dynamic	
Related Commands		
Note		

pool mc-buffer

pool <pool-name> **mc-buffer** <buffer> **reserved** <reserved> **shared** {alpha | max}
 <shared>
no pool <pool-name>

Configures pool.
 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.sp14
	reserved	The amount of shared memory for this buffer
	shared	The amount of shared memory for this buffer - When working in alpha mode, alpha can have the values 0, 1/128, 1/64 ... 1, 2, 4, ... 64, inf - When working in max mode, the shared size is defined as a percentage from the pool size
Default	N/A	
Configuration Mode	Config Interface Ethernet Egress Buffer	
History	3.6.1002	
Role	admin	
Example	switch (config interface ethernet 1/1 egress-buffer ePort.tc4)# pool iPool1 mc-buffer mx.sp0 reserved 90K shared alpha 1/8	
Related Commands		
Note		

pool description

pool <pool-name> description <description>
no pool <pool-name>

Configures the buffer description of a specific pool-name.
The no form of the command resets the values to their default.

Syntax Description	description	String text
Default	""	
Configuration Mode	Config Interface Ethernet Egress Buffer Config Interface Ethernet Ingress Buffer	
History	3.6.1002	
Role	admin	
Example	switch (config interface ethernet 1/1 ingress-buffer iPort.pg1)# pool iPool1 description myDescription	
Related Commands		
Note		

show buffers status

show buffers status interfaces ethernet <slot>/<port>

Displays buffer status

Syntax Description	<slot>/<port>	Ethernet interface				
Default	N/A					
Configuration Mode	Config					
History	3.6.1002					
Role	admin					
Example	switch (config)# show buffers status 1/25					
	Interface	Buffer	Resv [Byte]	Shared [%/a]	Usage [Byte]	MaxUsage [Byte]
	-----	-----	----	-----	-----	-----
	Eth1/25	iPort	192	1/128	0	0
		iPort	0	0	0	0
		iPort	0	0	0	0
		iPort	0	0	0	0
		iPort.pg0	0	0	0	0
		iPort.pg1	0	0	0	0
		iPort.pg2	0	0	0	0
		iPort.pg3	0	0	0	0
		iPort.pg4	0	0	0	0
		iPort.pg5	0	0	0	0
		iPort.pg6	0	0	0	0
		iPort.pg7	0	0	0	0
		iPort.pg9	19.5K	inf	0	0
		ePort	0	inf	0	0
		ePort	0	inf	0	0
		ePort	0	inf	0	0
		ePort	0	inf	0	0
		ePort.tc0	1.5K	2	0	0
		ePort.tc1	1.5K	2	0	0
		ePort.tc2	1.5K	2	0	0
		ePort.tc3	1.5K	2	0	0
		ePort.tc4	1.5K	2	0	0
		ePort.tc5	1.5K	2	0	0
		ePort.tc6	1.5K	2	0	0
		ePort.tc7	1.5K	2	0	0
		ePort.tc8	0	0	0	0
		ePort.tc9	0	0	0	0
ePort.tc10		0	0	0	0	
ePort.tc11		0	0	0	0	
ePort.tc12		0	0	0	0	
ePort.tc13	0	0	0	0		
ePort.tc14	0	0	0	0		
ePort.tc15	0	0	0	0		
ePort.tc16	96	inf	0	0		

Related Commands

Note

show buffers details

show buffers details interfaces ethernet <slot>/<port>

Displays buffer status in details.

Syntax Description	<slot>/<port>	Ethernet interface
Default	N/A	
Configuration Mode	Config	
History	3.6.1002	
Role	admin	

Example

```
switch (config)# show buffers details interfaces ethernet 1/25
Flags: Y - Lossy, L - Lossless
       S - Static, D - Dynamic
Shared size is in Bytes for static pool and in alphas for dynamic pool.
```

Interface: Eth1/25

Buffer	Resv [Byte]	Xoff [Byte]	Xon [Byte]	Shared [%/a]	Pool	Description
-----	-----	-----	-----	-----	-----	-----
iPort(Y)	192	-	-	1/128	iPool0(D)	
iPort(Y)	0	-	-	0	iPool1(D)	
iPort(Y)	0	-	-	0	iPool2(D)	
iPort(Y)	0	-	-	0	iPool3(D)	
iPort.pg0(Y)	0	-	-	0	iPool0(D)	Data
iPort.pg1(Y)	0	-	-	0	iPool0(D)	
iPort.pg2(Y)	0	-	-	0	iPool0(D)	
...						
iPort.pg7(Y)	0	-	-	0	iPool0(D)	
iPort.pg9(Y)	19.5K	-	-	inf	iPool0(D)	Control
ePort	0	-	-	inf	ePool0(D)	
ePort	0	-	-	inf	ePool1(D)	
ePort	0	-	-	inf	ePool2(D)	
ePort	0	-	-	inf	ePool3(D)	
ePort.tc0	1.5K	-	-	2	ePool0(D)	
ePort.tc1	1.5K	-	-	2	ePool0(D)	
ePort.tc2	1.5K	-	-	2	ePool0(D)	
...						
ePort.tc6	1.5K	-	-	2	ePool0(D)	
ePort.tc7	1.5K	-	-	2	ePool0(D)	
ePort.tc8	0	-	-	0	ePool0(D)	
ePort.tc9	0	-	-	0	ePool0(D)	
...						
ePort.tc15	0	-	-	0	ePool0(D)	
ePort.tc16	96	-	-	inf	ePool0(D)	Control
Switch-priority	Buffer					
-----	-----					
0	iPort.pg0					
1	iPort.pg0					
2	iPort.pg0					
3	iPort.pg0					
4	iPort.pg0					
...						
10	iPort.pg0					
11	iPort.pg0					
12	iPort.pg0					
13	iPort.pg0					
14	iPort.pg0					

Related Commands

Note

6 IP Routing

6.1 General

6.1.1 IP Interfaces

MLNX-OS supports the following 3 types of IP interfaces:

- VLAN interface
- Loopback interface
- Router port interface



Router port interfaces are not supported on SX10xx-xxxR and SX60xx-xxxR systems.

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. The user may configure up to 64 VLAN interfaces.

Each interface VLAN has the following attributes:

- Admin state
- Operational state
- MAC address
- IP address and mask
- MTU
- Description
- Set of counters

Loopback interface is a logical software entity where traffic transmitted to this interface is immediately received on the sending end.

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.

6.1.1.1 Configuring a VLAN Interface

➤ **To configure a VLAN interface:**

Step 1. Create a VLAN. Run:

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

Step 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
```

Step 3. There must be at least one interface in the operational state “UP”.

```
switch (config)# show interface ethernet 1/1 status
```

Port	Operational state	Speed	Negotiation
----	-----	-----	-----
Eth1/1	Up	40 Gbps	No-Negotiation

Step 4. Create a VLAN interface that matches the VLAN. Run:

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

Step 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
```

Step 6. Verify VLAN interface configuration. Run:

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

```
Vlan 10
Admin state: Enabled
Operational state: UP
Mac Address: 00:02:c9:5d:e0:f0
Internet Address: 10.10.10.10/24
Broadcast address: 10.10.10.255
MTU: 1500 bytes
Description: my-ip-interface
Counters: disabled
```

6.1.1.2 Configuring a Loopback Interface

➤ *To configure a loopback interface:*

Step 1. Create a loopback interface. Run:

```
switch (config)# interface loopback 2
switch (config interface loopback 2)#
```

Step 2. Configure an IP address on the loopback interface. Run:

```
switch (config interface loopback 2)# ip address 20.20.20.20 /32
```


Step 3. Verify loopback interface configuration. Run:

```
switch (config interface loopback 2)# show interfaces loopback 2

Loopback 2
  Internet Address: 20.20.20.20/32
  Broadcast address: 20.20.20.20
  MTU: 1500 bytes
  Description: my-loopback
switch (config) #
```

6.1.1.3 Configuring a Router Port Interface

Step 1. Enter an Ethernet interface's configuration context. Run:

```
switch (config)# interface ethernet 1/10
switch (config interface ethernet 1/10)#
```

Step 2. Configure the Ethernet interface to become a router port interface. Run:

```
switch (config interface ethernet 1/10)# no switchport force
```

Step 3. Configure an IP address on the router port interface. Run:

```
switch (config interface ethernet 1/10)# ip address 100.100.100.100 /24
```

Step 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
  Description: N/A
  Mac address: 00:02:c9:96:c6:d8
  MTU: 1500 bytes(Maximum packet size 1522 bytes)
  Flow-control: receive off send off
  Actual speed: 40 Gbps
  Width reduction mode: disabled
  DHCP client: Disabled
  IP Address: 100.100.100.100 /24
  Broadcast address: 100.100.100.255
  Arp timeout: 1500 seconds
  VRF: default
  MAC learning mode: Enabled
  Last clearing of "show interface" counters : 00:00:01
  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

```


6.1.2 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 Figure 37, 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

Figure 37: ECMP



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.

6.1.2.1 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.

Figure 38: Multiple Hash Functions



6.1.3 Virtual Routing and Forwarding



Only static IPv4 and ECMP are supported with VRF.

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.

6.1.4 Commands

6.1.4.1 General

ip l3

ip l3 [force]
no ip l3 [force]

Enables IP routing capabilities.
The no form of the command disables IP routing and removes its configuration.

Syntax Description	N/A
Default	If operating with Ethernet system profile: L3 If operating with VPI system profile: L2
Configuration Mode	Config
History	3.4.1802
Role	admin
Example	switch (config) # ip l3 force switch (config) #
Related Commands	N/A
Note	

vrf definition

vrf definition <vrf-name>

Creates the VRF.

Syntax Description	vrf-name	VRF session name
Default	N/A	
Configuration Mode	Config	
History	3.4.2008	
Role	admin	
Example	switch (config) # vrf definition my-vrf switch (config vrf definition my-vrf) #	
Related Commands	N/A	
Notes	Only 1 VRF is supported aside from the default VRF	

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	
Role	admin	
Example	switch (config) # routing-context vrf my-vrf switch (config) #	
Related Commands	N/A	
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	

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
Role	admin	
Example	<pre>switch (config) # ip routing vrf my-vrf switch (config) #</pre>	
Related Commands	N/A	
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.	

description

description <description>
no description force

Creates 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	
Role	admin	
Example	<pre>switch (config vrf definition my-vrf) # description vrf-description switch (config vrf definition my-vrf) #</pre>	
Related Commands	N/A	
Notes		

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	
Role	admin	
Example	<pre>switch (config vrf definition my-vrf) # rd 10.10.10.10:2 switch (config vrf definition my-vrf) #</pre>	
Related Commands	N/A	
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	

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	
Role	admin	
Example	switch (config interface ethernet 1/2) # vrf forwarding my-vrf switch (config interface ethernet 1/2) #	
Related Commands	N/A	
Notes		

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
Role	admin	
Example	<pre>switch (config) # show ip routing vrf all VRF Name: my-vrf ----- IP routing: disabled VRF Name: default ----- IP routing: enabled switch (config) #</pre>	
Related Commands	N/A	
Notes	If no routing-context is specified, the “routing-context” VRF is automatically displayed.	

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
Role	admin
Example	switch (config) # show routing-context vrf VRF active context: my-vrf switch (config) #
Related Commands	N/A
Notes	

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	
Role	admin	
Example	<pre>switch (config) # show vrf my-vrf VRF Info Name: my-vrf RD: 10.10.10.10:2 Description: Test VRF IP routing state: Enabled Protocols: IPv4 Interfaces: Eth1/2 switch (config) #</pre>	
Related Commands	N/A	
Notes	If no routing-context is specified, the “routing-context” VRF is automatically displayed.	

6.1.4.2 IP Interfaces

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	
Role	admin	
Example	switch (config interface ethernet 1/10)# no switchport force	
Related Commands		
Note		

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	
Role	admin	
Example	switch (config interface ethernet 1/10)# encapsulation dot1q vlan 10	
Related Commands		
Note		

6.1.4.3 Interface VLAN

interface vlan

interface vlan <vlan-id>
no interface vlan <vlan-id>

Creates a VLAN interface and enters the interface VLAN configuration mode.

The no form of the command deletes the VLAN interface.

Syntax Description	vlan-id A numeric range of 1-4094
Default	N/A
Configuration Mode	Config
History	3.2.0230
Role	admin
Example	<pre>switch (config) # interface vlan 10 switch (config interface vlan 10) #</pre>
Related Commands	<pre>ip routing vlan <vlan-id> switchport mode switchport access show interfaces vlan</pre>
Note	- 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 “switchport ...” - At least one interface belong to that VLAN must be in UP state

ip address

ip address <ip-address> <mask>
no ip address <ip-address> <mask>

Enters user-defined description for the interface.

Syntax Description	ip-address	IPv4 address
	mask	There are two possible ways to the mask: • /length (i.e. /24) • Network address (i.e. 255.255.255.0)
Default	0.0.0.0/0	
Configuration Mode	Config Interface VLAN	
History	3.2.0230	
Role	admin	
Example	switch (config interface vlan 10) # ip address 10.10.10.10 /24 switch (config interface vlan 10) #	
Related Commands	interface vlan show interfaces vlan	
Note		

ip address dhcp

ip address dhcp
no ip address dhcp

Enables DHCP on this VLAN interface.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config Interface VLAN
History	3.4.2008
Role	admin
Example	switch (config interface vlan 10) # ip address dhcp switch (config interface vlan 10) #
Related Commands	interface vlan show interfaces vlan
Note	

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	counters are disabled.
Configuration Mode	Config Interface VLAN
History	3.2.0230
Role	admin
Example	switch (config interface vlan 10) # counters switch (config interface vlan 10) #
Related Commands	counters interface vlan show interfaces vlan
Note	- Enabling counters for the router interface adds delay to the traffic stream - There are maximum of 16 counter sets

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	
Role	admin	
Example	switch (config interface vlan 10) # description my-ip-interface switch (config interface vlan 10) #	
Related Commands	interface vlan show interfaces vlan	
Note		

mtu

mtu <size> [force]
no mtu

Sets the MTU for the interface.
 The no form of the command sets the MTU to default.

Syntax Description	size	1500-9216.
	force	Forces command implementation.
Default	1522	
Configuration Mode	Config Interface VLAN	
History	3.2.0230	
Role	admin	
Example	switch (config interface vlan 10)# mtu 9216 switch (config interface vlan 10 #	
Related Commands	interface vlan show interfaces vlan	
Note		

shutdown

shutdown
no shutdown

Disables the interface.
The no form of the command enables the interface.

Syntax Description	N/A
Default	The interface is enabled.
Configuration Mode	Config Interface VLAN
History	3.1.0000
Role	admin
Example	<pre>switch (config interface vlan 20) # shutdown switch (config interface vlan 20) #</pre>
Related Commands	interface vlan
Note	

clear counters

clear counters

Clears the interface counters.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config Interface VLAN
History	3.2.0230
Role	admin
Example	switch (config interface vlan 10) # clear counters switch (config interface vlan 10) #
Related Commands	interface vlan counters
Note	

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
Role	admin
Example	switch (config interface vlan 10) # no ip icmp redirect
Related Commands	interface vlan counters
Note	ICMP redirect transmits messages to hosts alerting them about the existence of more efficient routes to a specific destination

show ip interface

show ip interface [vrf <vrf-name> | all] [brief]

Displays IP interfaces information per VRF.

Syntax Description	all		Displays information on all VRFs			
	brief		Displays IP interfaces information in a shortened form			
Default	N/A					
Configuration Mode	Any Command Mode					
History	3.4.2008					
Role	admin					
Example	<pre>switch (config) # show ip interface vrf all brief Interface Address/Mask Admin-state Oper-state MTU VRF mgmt0 10.224.22.27/24 Enabled Up 1500 default mgmt1 0.0.0.0/0 Enabled Down 1500 default Vlan 20 20.20.20.1/24 Enabled Down 1500 my-vrf Eth1/1 1.1.1.1/24 Enabled Down 1500 my-vrf Loopback 10 10.10.10.1/32 Enabled Up 1500 my-vrf Vlan 30 30.30.30.1/24 Enabled Down 1500 default Eth1/2 2.2.2.2/24 Enabled Down 1500 default Loopback 11 11.11.11.1/32 Enabled Up 1500 default switch (config) # show ip interface vrf my-vrf brief Interface Address/Mask Admin-state Oper-state MTU VRF Vlan 20 20.20.20.1/24 Enabled Down 1500 my-vrf Eth1/1 1.1.1.1/24 Enabled Down 1500 my-vrf Loopback 10 10.10.10.1/32 Enabled Up 1500 my-vrf switch (config) # show ip interface vrf default brief Interface Address/Mask Admin-state Oper-state MTU VRF mgmt0 10.224.22.27/24 Enabled Up 1500 default mgmt1 0.0.0.0/0 Enabled Down 1500 default Vlan 30 30.30.30.1/24 Enabled Down 1500 default Eth1/2 2.2.2.2/24 Enabled Down 1500 default Loopback 11 11.11.11.1/32 Enabled Up 1500 default switch (config) #</pre>					
Related Commands	N/A					
Notes	If no routing-context is specified, the “routing-context” VRF is automatically displayed.					

6.1.4.4 Loopback Interface

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	A numeric range of 0-31
Default	N/A	
Configuration Mode	Config	
History	3.2.3000	
Role	admin	
Example	switch (config) # interface loopback 10 switch (config interface loopback 10) #	
Related Commands		
Note	- 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	

ip address

ip address <ip-address> <mask>
no ip address <ip-address> <mask>

Enters user-defined description for the interface.

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)
Default	0.0.0.0/0	
Configuration Mode	Config Interface Loopback	
History	3.3.5006	
Role	admin	
Example	switch (config interface loopback 10) # ip address 10.10.10.10 /32	
Related Commands	interface loopback	
Note		

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 Loopback	
History	3.3.5006	
Role	admin	
Example	switch (config interface loopback 10) # description my-ip-interface	
Related Commands	interface loopback	
Note		

show interfaces loopback

show interface loopback <id>

Shows the attribute of the interface loopback.

Syntax Description	id	A numeric range of 1-32
Default	N/A	
Configuration Mode	Config	
History	3.2.3000	
Role	admin	
Example	switch (config) # show interfaces loopback 2 Loopback 2 Internet Address: 2.2.2.2/32 Broadcast address: 2.2.2.2 MTU: 1500 bytes Description: my-loopback switch (config) #	
Related Commands		
Note		

6.1.4.5 Routing and ECMP

ip route

ip route [vrf <vrf-name>] <IP prefix> <netmask> <next hop IP address>
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 the mask: • /length (i.e. /24) • Network address (i.e. 255.255.255.0)
	next hop IP address	IP address of the next hop.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
	3.4.2008	Added VRF parameter
Role	admin	
Example	switch (config) # ip route vrf my-vrf 80.80.80.0 /24 20.20.20.2	
Related Commands	N/A	
Notes	If no routing-context is specified, the “routing-context” VRF is automatically configured.	

ip load-sharing

ip load-sharing <type>
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 • traffic-class – traffic class • flow-label – flow label • all – all options
Default	all	
Configuration Mode	Config	
History	3.2.0230	
	3.5.1000	Added flow-label parameter and updated Note section
Role	admin	
Example	<pre>switch (config) # ip load-sharing all switch (config) # show ip load-sharing Load sharing: all switch (config)</pre>	
Related Commands	ip route	
Note	The parameter “traffic-class” is available on SwitchX® based systems only	

show ip route

show ip route [vrf [<vrf-name> | all]] [-a | static | summary]

Displays routing table of VRF instance.

Syntax Description	all	Displays routing tables for all VRF instances
	-a	Displays static routes currently inactive due to the interface being down
	static	Displays static route
	summary	Displays route summary
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.1.0000	First version
	3.3.3500	Added Distance/Metric column
	3.4.0000	Added -a parameter
	3.4.2008	Added VRF parameter
	3.4.3000	Updated Notes section
Role	admin	

Example

```
switch (config) # show ip route vrf my-vrf
```

```
VRF Name:      my-vrf
```

Destination	Mask	Gateway	Interface	Source	Distance/Metric
10.10.10.1	255.255.255.255	0.0.0.0	loopback10	direct	0/0
20.20.20.0	255.255.255.0	0.0.0.0	vlan20	direct	0/0
80.80.80.0	255.255.255.0	20.20.20.2	vlan20	static	1/0

```
switch (config) # show ip route vrf my-vrf static
```

```
VRF Name:      my-vrf
```

Destination	Mask	Gateway	Interface	Source	Distance/Metric
80.80.80.0	255.255.255.0	20.20.20.2	vlan20	static	1/0

```
switch (config) # show ip route vrf my-vrf summary
```

```
VRF Name:      my-vrf
```

Route Source	Routes
direct	2
static	1
ospf	0
bgp	0
DHCP	0
Total	3

```
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

```
switch (config) #
```

Related Commands

ip route

Notes

- If no routing-context is specified, the “routing-context” VRF is automatically displayed
- If no default route exists, then the message “Route not found” is printed

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	3.2.0230
Role	admin
Example	<pre>switch (config) # show ip load-sharing Load sharing: all switch (config) #</pre>
Related Commands	ip load-sharing
Note	

6.1.4.6 Network to Media Resolution (ARP)

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	
Role	admin	
Example	switch (config) # ip arp vrf my-vrf 20.20.20.2 aa:bb:cc:dd:ee:ff	
Related Commands	N/A	
Notes	If no routing-context is specified, the “routing-context” VRF is automatically configured.	

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 (in seconds) that an entry remains in the ARP cache. Range: 240-28800.
Default	1500 seconds	
Configuration Mode	Config Interface Ethernet Config Interface VLAN Config Interface Port Channel	
History	3.2.0230	
	3.5.1000	Updated Note section
Role	admin	
Example	<pre>switch (config) # ip arp timeout 2000 switch (config) # show ip arp ARP Timeout: 2000 Total number of entries: 55 IP Address MAC Address Interface 1.0.0.2 00:02:c9:5c:30:40 Vlan11 1.0.0.3 00:11:22:33:44:55 Vlan11 2.0.0.2 00:02:c9:5c:30:40 Vlan12 3.0.0.2 00:02:c9:5c:30:40 Vlan13 4.0.0.2 00:02:c9:5c:30:40 Vlan14 switch (config) #</pre>	
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	

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	
History	3.4.2008	Added VRF parameter
Role	admin	
Example	switch (config) # clear ip arp vrf my-vrf switch (config) #	
Related Commands	ip arp show ip arp	
Notes	If no routing-context is specified, the “routing-context” VRF is automatically configured.	

show ip arp

show ip arp [vrf [<vrf-name> | all]] [interface <type> | count]

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
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.3000	
	3.4.2008	Added VRF parameter
Role	admin	
Example	<pre>switch (config) # show ip arp vrf my-vrf VRF Name: my-vrf ----- Total number of entries: 2 Address Type Hardware Address Interface ----- 20.20.20.2 Static ETH AA:AA:AA:BB:BB:BB vlan 20 1.1.1.2 Static ETH 00:11:22:33:44:55 eth 1/1 switch (config) # show ip arp vrf my-vrf interface ethernet 1/1 VRF Name: my-vrf ----- Total number of entries: 1 Address Type Hardware Address Interface ----- 1.1.1.2 Static ETH 00:11:22:33:44:55 eth 1/1 switch (config) # show ip arp vrf my-vrf interface vlan 20 VRF Name: mmm ----- Total number of entries: 1 Address Type Hardware Address Interface ----- 20.20.20.2 Static ETH AA:AA:AA:BB:BB:BB vlan 20 switch (config) #</pre>	

Related Commands	ip arp

Notes	If no routing-context is specified, the “routing-context” VRF is automatically displayed.

6.1.4.7 IP Diagnostic Tools

ping

ping [vrf <vrf-name>] [-LRUbdnqrvVaA] [-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		
	<table> <tr> <td>vrf</td><td>Specifies VRF instance name</td></tr> </table>	vrf	Specifies VRF instance name
vrf	Specifies VRF instance name		
Default	N/A		
Configuration Mode	Config		
History	3.1.0000 3.4.2008 Added VRF parameter		
Role	admin		
Example	<pre>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 switch (config) #</pre>		
Related Commands	traceroute		
Note	When using -I option use the interface name + interface number, for example "ping -I vlan10"		

traceroute

traceroute [vrf <vrf-name>] [-4dFITUnrAV] [-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 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.
-q	Sets the number of probes per each hop. Default is 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
Role	admin
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 switch (config) #
Related Commands	
Note	- The following flags are not supported: -6, -l, -A - When using -i option use the interface name + interface number, for example “traceroute -i vlan10”

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
Role	admin	
Example	<pre>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 switch (config) #</pre>	
Related Commands	N/A	
Note	- 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.	

6.1.4.8 QoS

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
Role	admin
Example	switch (config) # qos map dscp-to-pcp preserve-pcp switch (config) #
Related Commands	
Note	

6.2 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 to 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 Mellanox L2 VMS, please refer to [Mellanox Virtual Modular Switch Reference Guide](#).

6.2.1 Router ID

The router ID is a 32-bit number assigned to the router running the OSPF protocol. This number uniquely identifies the router within an Autonomous System.

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 loopback IP address;
- Otherwise, the lowest IP address is elected as router ID

6.2.2 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 cost, 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.

6.2.3 Configuring OSPF

Figure 39: OSPF Basic Topology



Precondition steps:



The following configuration example refers to Router 2 in [Figure 39](#). 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`.

Step 1. Make sure an L3 license is installed. For a list of the available licenses see [Section 2.4, “Licenses,”](#) on page 63.

Step 2. Enable IP routing functionality. Run:.

```
switch (config)# ip routing
```

Step 3. Enable the desired VLAN. Run:.

```
switch (config)# vlan 10
switch (config)# vlan 20
```

Step 4. 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
```

Step 5. Create a VLAN interface. Run:

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

Step 6. Apply IP address to the VLAN interface. Run:

```
switch (config interface vlan 10)# ip address 10.10.10.2 /16
```

Step 7. Enable the interface. Run:

```
switch (config interface vlan 10)# no shutdown
```

Step 8. Create a second VLAN interface. Run:

```
switch (config)# interface vlan 20
```

Step 9. Apply IP address to the second VLAN interface. Run:

```
switch (config interface vlan 20)# ip address 10.10.20.2 /16
```

Step 10. Enable the second interface. Run:

```
switch (config interface vlan 20)# no shutdown
```

Basic OSPF Configuration:

Step 1. To enable OSPF configuration run:

```
switch (config)# protocol ospf
```

Step 2. To create a router OSPF instance run:

```
switch (config)# router ospf
```



Only one instance of OSPF is supported.

Step 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
```

6.2.4 Verifying OSPF

➤ **To verify OSPF configuration and status:**

Step 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 40Gb
Initial SPF schedule delay 1 msec
SPF Hold time 10 msec
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

switch (config) #
```

- Step 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

switch (config) #
```

Step 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
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
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

switch (config) #
```

6.2.5 Commands

6.2.5.1 Config

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	OSPF feature is disabled.
Configuration Mode	Config
History	3.3.3500
Role	admin
Example	switch (config)# protocol ospf
Related Commands	ip routing
Note	

router ospf

router ospf [<process-id> [vrf <vrf-name>]]
no router ospf [<process-id> [vrf <vrf-name>]]

Enters router OSPF configuration mode, and creates default OSPF instance on specific VRF with specific Process ID if one does not exist. The no form of the command deletes the OSPF instance.

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
Role	admin	
Example	switch (config)# router ospf 2 vrf myvrf switch (config router ospf 2)#	
Related Commands	N/A	
Note	Only one OSPF instance is supported	

6.2.5.2 Config Router

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 Autonomous System. 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 loopback IP address; • Otherwise, the lowest IP address is elected as router ID.	
Configuration Mode	Config OSPF Router	
History	3.3.3500	
Role	admin	
Example	switch (config router ospf)# router-id 10.10.10.10	
Related Commands	N/A	
Note		

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
Role	admin
Example	switch (config router ospf)# shutdown
Related Commands	N/A
Note	

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 Gbps (default if not specified)
	Mbps	Value in Mbps
Default	40 Gbps	
Configuration Mode	Config OSPF Router	
History	3.3.3500	
Role	admin	
Example	switch (config router ospf)# auto-cost reference-bandwidth 10 Gbps	
Related Commands	N/A	
Note		

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	
Role	admin	
Example	switch (config router ospf)# distance 100	
Related Commands	N/A	
Note		

redistribute

redistribute {bgp | direct | static}
no redistribute {bgp | direct | static}

Import 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	Redistribute directly connected routes.
	bgp	Redistribute routes from BGP protocol.
	static	Redistribute static configured routes.
Default	Disable (no redistribution)	
Configuration Mode	Config OSPF Router	
History	3.2.1000	
Role	admin	
Example	switch (config router ospf)# redistribute direct	
Related Commands	N/A	
Note	Routes from multiple protocols can be imported in parallel.	

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 is 0-100 milliseconds.
	spf-hold	The minimum delay between two consecutive delay calculations. Range is 0-1000 milliseconds.
Default	spf-delay: 1 millisecond spf-hold: 10 millisecond	
Configuration Mode	Config OSPF Router	
History	3.3.3500	
Role	admin	
Example	switch (config router ospf)# timers throttle spf 100 1000	
Related Commands	N/A	
Note		

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 is 0-4294967295.
	cost	The cost for the default summary route. Range is 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	
Role	admin	
Example	switch (config router ospf)# area 0 default-cost 100	
Related Commands	N/A	
Note	Base cost for all calculation is 56GbE.	

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 is 0-4294967295.
	ip-address	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	Disabled	
Configuration Mode	Config OSPF Router	
History	3.3.3500	
Role	admin	
Example	switch (config router ospf)# area 0 range 10.10.10.10 /24	
Related Commands	N/A	
Note		

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 is 0-4294967295.
	no-summary	Summary route will not be advertised into the stub area.
Default	Summary route will be advertised.	
Configuration Mode	Config OSPF Router	
History	3.3.3500	
Role	admin	
Example	switch (config router ospf)# area 0 stub	
Related Commands	N/A	
Note		

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 is 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 is 1-2.
	m-value	Metric value for OSPF. Range is 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	
Role	admin	
Example	switch (config router ospf)# area 0 nssa	
Related Commands	N/A	
Note	An area can be either stub, NSSA or normal.	

no area

no area <area-id>

Deletes OSPF area and its related configuration.

Syntax Description	area-id	OSPF area ID Range is 0-4294967295
Default	N/A	
Configuration Mode	Config OSPF Router	
History	3.3.3500	
Role	admin	
Example	switch (config router ospf)# no area 1	
Related Commands	N/A	
Note	The command fails if the area is attached to active interfaces.	

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	
Role	admin	
Example	switch (config router ospf)# summary-address 10.10.10.10 /24	
Related Commands	N/A	
Note	Maximum of 1500 summarized IP addresses can be configured.	

6.2.5.3 Interface

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	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	
Role	admin	
Example	switch (config interface vlan 10)# ip ospf cost 100	
Related Commands	N/A	
Note		

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, in seconds. Range is 1-65535.
Default	40	
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	
Role	admin	
Example	switch (config interface vlan 10)# ip ospf dean-interval 10	
Related Commands	N/A	
Note	The value must be the same for all nodes on the network.	

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, in seconds. Range is 1-65535.
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	
Role	admin	
Example	switch (config interface vlan 10)# ip ospf hello-interval 20	
Related Commands	N/A	
Note	The value must be the same for all nodes on the network.	

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 is 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	
Role	admin	
Example	switch (config interface vlan 10)# ip ospf priority 100	
Related Commands	N/A	
Note	- 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.	

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. The options are 'broadcast' or 'point-to-point'.
Default	broadcast	
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	
Role	admin	
Example	switch (config interface vlan 10)# ip ospf network point-to-point	
Related Commands	N/A	
Note	- The network type influences the behavior of the OSPF interface. An OSPF network type is usually broadcast, which uses OSPF multicasting 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 should have the same network type.	

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 in seconds. Range is 0-3600.
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	
Role	admin	
Example	switch (config interface vlan 10)# ip ospf retransmit-interval 10	
Related Commands	N/A	
Note		

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
Role	admin
Example	switch (config interface vlan 10)# ip ospf passive-interface
Related Commands	N/A
Note	

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 is 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	
Role	admin	
Example	switch (config interface vlan 10)# ip ospf transmit-delay 2	
Related Commands	N/A	
Note		

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
Role	admin
Example	switch (config interface vlan 10)# ip ospf shutdown
Related Commands	N/A
Note	

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 (no)	
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	
Role	admin	
Example	switch (config interface vlan 10)# ip ospf authentication	
Related Commands	N/A	
Note	- Without message-digest option, a simple password authentication will be used. - Message-digest authentication can be enabled only if a key is configured.	

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	
Role	admin	
Example	switch (config interface vlan 10)# ip ospf authentication-key 0 mycleartextpassword	
Related Commands	N/A	
Note	- 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.	

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 (no)	
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	
Role	admin	
Example	switch (config interface vlan 10)# ip ospf message-digest-key mykeyid md5 7 mykey	
Related Commands	N/A	
Note	The user cannot delete the last key until authentication is disabled.	

ip ospf area

ip ospf area <area-id>
no ip ospf area

Sets 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 is 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	
Role	admin	
Example	switch (config interface vlan 10)# ip ospf area 0	
Related Commands	N/A	
Note		

6.2.5.4 Show

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
Role	admin	
Example	<pre>switch (config)# show ip ospf 2 vrf myvrf</pre> 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 msec SPF Hold time 5000 msec 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	N/A	
Note		

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
Role	admin	
Example	<pre>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</pre>	
Related Commands	N/A	
Note		

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 is 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
Role	admin	
Example	<pre>switch (config)# show ip ospf database 2 vrf myvrf</pre> OSPF Router with ID (2.2.2.2) (Process ID 2 VRF myvrf) <pre> 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 </pre>	

Related Commands	N/A

Note

show ip ospf interface

show ip ospf interface [<process-id>] [vlan <vlan-id>] [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																					
Role	admin																						
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 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 <table><thead><tr><th>Interface Id</th><th>Area</th><th>Cost</th><th>State</th><th>Neighbors</th><th>Status</th></tr></thead><tbody><tr><td>Vlan21</td><td>0.0.0.0</td><td>10</td><td>Enabled</td><td>1</td><td>Up</td></tr><tr><td>Ethernet1/22</td><td>0.0.0.0</td><td>1</td><td>Enabled</td><td>1</td><td>Up</td></tr></tbody></table>					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
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	N/A

Note

show ip ospf neighbors

show ip ospf neighbors [vlan <vlan-id>] [<neighbor-id>] [vrf <vrf-name>]

Displays the OSPF related interface neighbor configuration.

Syntax Description	vlan <vlan-id>	Displays OSPF interface configuration and status per VLAN interface
	neighbor-id	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
Role	admin	
Example	<pre>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 Interface Ethernet 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 interface ethernet 1/22 vrf myvrf Neighbor 1.1.1.1, interface address 22.22.22.1 In the area 0.0.0.0 via interface Ethernet 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</pre>	

Related Commands	N/A

Note

show ip ospf request-list

show ip ospf request-list <neighbor-id> vlan <vlan-id>

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.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.3500	
Role	admin	
Example	<pre>Router# show ip ospf request-list 40.40.40 ethernet 2/1 OSPF Process ID p1 Neighbor 40.40.40.40, interface Ethernet2/1, address 192.0.2.1 1 LSAs on request-list Type LS ID ADV RTR Seq NO Age Checksum 1 192.0.2.12 192.0.2.12 0x8000020D 8 0x6572</pre>	
Related Commands	N/A	
Note		

show ip ospf retransmission-list

show ip ospf retransmission-list <neighbor-id> vlan <vlan-id>

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.
	vlan-id	Filers the output per a specific VLAN ID.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.3500	
Role	admin	
Example	<pre>Router# show ip ospf retransmission-list 192.0.2.11 ethernet 2/1 OSPF Router with ID (192.0.2.12) (Process ID 1) Neighbor 192.0.2.11, interface Ethernet2/1 address 209.165.201.11 Link state retransmission due in 3764 msec, Queue length 2 Type LS ID ADV RTR Seq NO Age Checksum 1 192.0.2.12 192.0.2.12 0x80000210 0 0xB196</pre>	
Related Commands	N/A	
Note		

show ip ospf summary-address

show ip ospf summary-address

Displays a list of all summary address redistribution information configured on the OSPF.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.3500
Role	admin
Example	<pre>switch (config)# show ip ospf summary-address Display of Summary addresses for External Routes and area ranges for the summary LSAs OSPF Process default OSPF External Summary Address and area-range Configuration Information ----- Network Mask Area Advertise LSA type Metric Tag ----- 1.1.1.1 255.255.255.0 NA Advertise Type5 10 0 2.2.2.0 255.255.255.0 10.10.10.10 Not Advertise Type3 10 0</pre>
Related Commands	N/A
Note	

6.3 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.


6.3.1 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.

6.3.2 Configuring BGP

Figure 40: Basic BGP Configuration



Follow these steps for basic BGP configuration on two switches (Router 1 and Router 2):

Preconditions:

Step 1. Make sure the license installed supports L3.

Step 2. Enable IP routing functionality. Run:

```
switch (config)# ip routing
```

Step 3. Enable the desired VLAN. Run:

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



The same VLAN must be configured on both switches.

Step 4. Add this VLAN to the desired interface. Run:

```
switch (config)# interface ethernet 1/1  
switch (config ethernet 1/1)# switchport access vlan 10
```

Step 5. Create a VLAN interface. Run:

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

Step 6. Apply IP address to the VLAN interface on Router 1. Run:

```
switch (config interface vlan 10)# ip address 10.10.10.1 /24
```

Step 7. Apply IP address to the VLAN interface on Router 2. Run:

```
switch (config interface vlan 10)# ip address 10.10.10.2 /24
```

Step 8. Enable the interface. Run:

```
switch (config interface vlan 10)# no shutdown
```

Configure BGP:

Step 1. Enable BGP. Run:

```
switch (config)# protocol bgp
```

Step 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 similar to the local AS number of the configured router.

Step 3. Configure BGP Router 1 neighbor. Run:

```
switch (config router bgp 100)# neighbor 10.10.10.2 remote-as 100
```

Step 4. Configure BGP Router 2 neighbor. Run:

```
switch (config router bgp 100)# neighbor 10.10.10.1 remote-as 100
```

6.3.3 Verifying BGP

Step 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.
- In case 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.

Step 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
switch (config)#
```

You should be able to see running BGP counters and ESTABLISHED state per active neighbor.

6.3.4 Commands

6.3.4.1 Config

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
Role	admin
Example	switch (config)# protocol bgp switch (config)#
Related Commands	ip routing
Note	

clear ip bgp

clear ip bgp [{<ip-address> | all} [soft] [in]]

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
Default	N/A	
Configuration Mode	Config	
History	3.3.5006	First release
	3.3.5200	Updated description
	3.6.3004	Removed “out” parameter
Role	admin	
Example	switch (config)# clear ip bgp all switch (config)#	
Related Commands	N/A	
Note	This command removes BGP IPv4 learned routes from the routing table, reads all routes from designated peers, and sends routes to those peers as required.	

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: 1-65535.
Default	N/A	
Configuration Mode	Config	
History	3.3.5006	First version
	3.3.5200	Updated syntax description
Role	admin	
Example	switch (config)# router bgp 100 switch (config) router bgp 100)#	
Related Commands	ip routing	
Note		

6.3.4.2 Config Router

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
Role	admin
Example	switch (config router bgp 100)# no shutdown
Related Commands	
Note	

aggregate-address

aggregate-address <prefix> [summary-only] [as-set] [attribute-map]
no aggregate-address <prefix> [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	prefix	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	
Role	admin	
Example	switch-e07c04 [standalone: master] (config router bgp 4) # aggregate-address 3.5.3.7 /32	
Related Commands		
Note	- 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	

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
Role	admin	
Example	switch (config router bgp 100)# bestpath as-path multipath-relax	
Related Commands	maximum-paths	
Note	- 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	

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
Role	admin
Example	switch (config router bgp 100)# bgp fast-external-fallover
Related Commands	maximum-paths
Note	Although this feature improves BGP conversion time, it may cause instability in your BGP table due to a flapping interface.

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	
Role	admin	
Example	switch (config router bgp 100)# bgp listen limit 101	
Related Commands		
Note		

bgp listen range peer-group

bgp listen range <ip-prefix> <length> 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-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	100	
Configuration Mode	Config Router BGP	
History	3.4.0000	
	3.6.3004	Added note
Role	admin	
Example	switch (config router bgp 100)# bgp listen range 10.10.10.10 /24 peer-group my-group remote-as 13	
Related Commands		
Note	- To create a static peer group, use the command <code>neighbor peer-group</code> - 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	

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	First version
	3.4.0000	Updated syntax description
	3.6.3004	Added “force” parameter
Role	admin	
Example	switch (config router bgp 100)# cluster-id 10.10.10.10	
Related Commands	N/A	
Note		

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
Role	admin
Example	switch (config router bgp 100)# client-to-client reflection
Related Commands	N/A
Note	

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: 200 internal: 200 local: 200	
Configuration Mode	Config Router BGP	
History	3.3.5006	
Role	admin	
Example	switch (config router bgp 100)# distance 10 20 30	
Related Commands	N/A	
Note	• 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.	

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	
Role	admin	
Example	switch (config router bgp 100)# graceful-restart stalepath-time 350	
Related Commands	N/A	
Note		

graceful-restart helper

graceful-restart helper
no graceful-restart helper

Enables BGP graceful restart helper mode on the switch for all BGP neighbors.

The no form of the command disables BGP graceful restart helper mode on the switch for all BGP neighbors.

Syntax Description	N/A
Default	Graceful restart is enabled
Configuration Mode	Config Router BGP
History	3.4.0000 3.6.3004 Updated Note section
Role	admin
Example	switch (config router bgp 100)# graceful-restart helper
Related Commands	N/A
Note	- 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 - This parameter can only be configured when BGP is admin-down state

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.
Default	1	
Configuration Mode	Config Router BGP	
History	3.3.5006	
	3.3.5200	Updated description and notes
Role	admin	
Example	<pre>switch (config router bgp 100)# maximum-paths ibgp 10 switch (config router bgp 100)#</pre>	
Related Commands	N/A	
Note	- 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.	

neighbor advertisement-interval

neighbor {<ip-address> | <peer-group-name>} advertisement-interval <delay>
no neighbor {<ip-address> | <peer-group-name>} 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	ip-address	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
Default	30 seconds	
Configuration Mode	Config Router BGP	
History	3.4.0000	First version
	3.6.3004	Updated description of “delay” parameter
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 advertisement-interval 90	
Related Commands		
Note		

neighbor allowas-in

neighbor {<ip-address> | <peer-group-name>} allowas-in [number]
no neighbor {<ip-address> | <peer-group-name>} 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
Default	N/A	
Configuration Mode	Config Router BGP	
History	3.4.0000	First version
	3.6.3004	Updated description of "number" parameter
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 allowas-in 2	
Related Commands	ip routing router bgp <as-number>	
Note	Neighbors from the same AS as the router are considered as iBGP peers, and neighbors from other ASs are considered eBGP peers.	

neighbor description

neighbor {<ip-address> | <peer-group-name>} description <string>
no neighbor {<ip-address> | <peer-group-name>} 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
Default	No description	
Configuration Mode	Config Router BGP	
History	3.3.5006	First version
	3.3.5200	Updated example
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 description The next door neighbor	
Related Commands	N/A	
Note	The peer description only appears in the show commands.	

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 disables connecting to external peers.

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	First version
	3.3.5200	Updated default
	3.6.3004	Updated description of “ttl” parameter
Role	admin	
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>	
Note	The command does not establish the multi-hop if the only route to the peer is the default route (0.0.0.0).	

neighbor export-localpref

neighbor {<ip-address> | <peer-group-name>} export-localpref <value>
no neighbor {<ip-address> | <peer-group-name>} 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
Default	100	
Configuration Mode	Config Router BGP	
History	3.4.0000	First version
	3.6.3004	Updated description of “value” parameter
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 export-localpref 100	
Related Commands		
Note		

neighbor import-localpref

neighbor {<ip-address> | <peer-group-name>} import-localpref <value>
no neighbor {<ip-address> | <peer-group-name>} import-localpref

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
Default	100	
Configuration Mode	Config Router BGP	
History	3.4.0000	First version
	3.6.3004	Updated description of “value” parameter
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 import-localpref 100	
Related Commands		
Note		

neighbor local-as

neighbor {<ip-address> | <peer-group-name>} local-as <as-id> [no-prepend | replace-as]

no neighbor {<ip-address> | <peer-group-name>} 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
	as-id	Range: 0-4294967295 where “12000” configures the default, and prevents the system from inheriting this parameter’s group configuration
	no-prepend	local-as number is not pre-pended to the routes received from external neighbors
	replace-as	Prepends only the local autonomous system number (as configured with the IP address argument) to the AS path attribute
Default	12000	
Configuration Mode	Config Router BGP	
History	3.4.0000	First version
	3.6.3004	Updated description of “as-id” parameter
Role	admin	
Example	switch-e07c04 [standalone: master] (config router bgp 4) # neighbor 100.100.100.100 local-as 123	
Related Commands	ip routing neighbor <ip-address> remote-as <as-number>	
Note	- 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 <code>clear ip bgp</code>	

neighbor maximum-prefix

neighbor {<ip-address> | <peer-group-name>} maximum-prefix <maximum> [warning-only]

no neighbor {<ip-address> | <peer-group-name>} 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
Default	12000	
Configuration Mode	Config Router BGP	
History	3.4.0000	First version
	3.6.3004	Updated description of “maximum” parameter
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 maximum-prefix 12000 warning-only	
Related Commands	ip routing neighbor <ip-address> remote-as <as-number>	
Note		

neighbor next-hop-peer

neighbor {<ip-address> | <peer-group-name>} next-hop-peer [disable]
no neighbor {<ip-address> | <peer-group-name>} next-hop-peer

Configures the switch to list the peer address as the next hop in routes that it receives from the specified peer BGP-speaking neighbor or members of the specified peer group.

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
Default	no next-hop-peer	
Configuration Mode	Config Router BGP	
History	3.3.5006	
	3.6.3004	Added "disable" parameter
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 next-hop-peer	
Related Commands		
Note	This command overrides the next hop for all routes received from this neighbor or peer group	

neighbor next-hop-self

neighbor {<ip-address> | <peer-group-name>} next-hop-self [disable]
no neighbor {<ip-address> | <peer-group-name>} 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
Default	no next-hop-self	
Configuration Mode	Config Router BGP	
History	3.3.5006	
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 next-hop-self	
Related Commands	neighbor <ip-address> remote-as <as-number>	
Note	- 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.	

neighbor password

**neighbor {<ip-address> | <peer-group-name>} password [<encryption>]
<string>
no neighbor {<ip-address> | <peer-group-name>} 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
Default	no neighbor password	
Configuration Mode	Config Router BGP	
History	3.4.0000	First version
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 password 7 admin123	
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	

neighbor no-password

neighbor {<ip-address> | <peer-group-name>} no-password

Disables authentication for peer without inheritance.

Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
Default	N/A	
Configuration Mode	Config Router BGP	
History	3.6.3004	First version
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 no-password	
Related Commands	neighbor password	
Note		

neighbor peer-group

1. neighbor {<ip-address>} peer-group <peer-group-name>
2. neighbor {<peer-group-name>} peer-group
3. no neighbor {<ip-address>} 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
Default	N/A	
Configuration Mode	Config Router BGP	
History	3.4.0000	First version
	3.6.3004	Added notes
Role	admin	
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

Note

- 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 retains 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
-

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	ip-address	A BGP peer IP address
	peer-group-name	Peer group name
	as-number	The BGP peer as-number. Range: 1-65535.
Default	N/A	
Configuration Mode	Config Router BGP	
History	3.3.5006	First version
	3.3.5200	Updated description and note
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 remote-as 200 switch (config router bgp 100)#	
Related Commands	ip routing router bgp <as-number>	
Note	Neighbors from the same AS as the router are considered as iBGP peers, and neighbors from other ASs are considered eBGP peers.	

neighbor remove-private-as

neighbor {<ip-address> | <peer-group-name>} remove-private-as [disable]
no neighbor {<ip-address> | <peer-group-name>} 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	ip-address	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
Default	N/A	
Configuration Mode	Config Router BGP	
History	3.4.0000	First version
Role	admin	
Example	<pre>switch (config router bgp 100)# neighbor 10.10.10.10 remove-private-as switch (config router bgp 100)#</pre>	
Related Commands	<pre>ip routing router bgp <as-number></pre>	
Note	• 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.	

neighbor route-map

```
neighbor {<ip-address> | <peer-group-name>} route-map <route-map-name>
[in | out]
no neighbor {<ip-address> | <peer-group-name>} route-map <route-map-name>
[in | out]
```

Configures a route map to inbound BGP routes.
The no form of the command undoes the configuration.

Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	route-map-name	String. The name of the route-map
	in	Applies route map to inbound routes
	out	Applies route map to out-bound routes
Default	N/A	
Configuration Mode	Config Router BGP	
History	3.3.5006	First version
	3.3.5200	Updated notes and default
	3.4.1100	Added “out” parameter
	3.6.3004	Added note
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 route-map MyRoute-Map in	
Related Commands	neighbor <ip-address> remote-as <as-number> route-map <map-name> [deny permit] [sequence-number] clear ip bgp {<ip-address> all}	
Note	- 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	

neighbor no-route-map

neighbor {<ip-address> | <peer-group-name>} no-route-map

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
Default	N/A	
Configuration Mode	Config Router BGP	
History	3.6.3004	
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 no-route-map	
Related Commands	neighbor <ip-address> remote-as <as-number> route-map <map-name> [deny permit] [sequence-number]	
Note		

neighbor route-reflector-client

neighbor {<ip-address> | <peer-group-name>} route-reflector-client [disable]
no neighbor {<ip-address> | <peer-group-name>} route-reflector-client

Sets the neighbor as a client but does not set up the reflection itself.
The no form of the command disables route reflection for the specific peer.

Syntax Description	ip-address	IP address of the neighbor.
	peer-group-name	Peer group name
	disable	Unsets neighbor as route reflector client and prevents the system from inheriting this parameter's group configuration
Default	N/A	
Configuration Mode	Config Router BGP	
History	3.3.5006	First version
	3.3.5200	Updated notes and default
	3.6.3004	Added “disable” parameter
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 route-reflector-client	
Related Commands		
Note		

neighbor send-community

neighbor {<ip-address> | <peer-group-name>} **send-community** [extended] [disable]

no neighbor {<ip-address> | <peer-group-name>} **send-community** [extended]

Configures the switch to send community attributes to the specified BGP neighbor.

The no form of the command disables sending community attributes for the specified peer.

Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	extended	Sends extended community attributes to neighbor
	disable	Disables sending community attributes for the specified peer and prevents the system from inheriting this parameter's group configuration
Default	Enabled	
Configuration Mode	Config Router BGP	
History	3.4.0000	First version
	3.6.3004	Added "disable" parameter
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 send-community	
Related Commands	N/A	
Note		

neighbor shutdown

neighbor {<ip-address> | <peer-group-name>} shutdown [disable]
no neighbor {<ip-address> | <peer-group-name>} 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
Default	Enabled	
Configuration Mode	Config Router BGP	
History	3.3.5006	First version
	3.3.5200	Updated note
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 shutdown	
Related Commands	N/A	
Note	- 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	

neighbor timers

neighbor {<ip-address> | <peer-group-name>} timers <keep-alive> <hold-time>
no neighbor {<ip-address> | <peer-group-name>} 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
Default	keep-alive: 60 seconds hold-time: 180 seconds	
Configuration Mode	Config Router BGP	
History	3.3.5006	First version
	3.3.5200	Updated description
	3.6.3004	Updated “hold-time” and “keep-alive” parameter’s syntax description
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 timers 65 195	
Related Commands	neighbor <ip-address> remote-as <as-number>	
Note	Hold time must be at least 3 seconds and should be three times longer than the keep-alive setting.	

neighbor transport connection-mode passive

neighbor {<ip-address> | <peer-group-name>} transport connection-mode passive [disable]
no neighbor {<ip-address> | <peer-group-name>} 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
Default	TCP sessions initiated	
Configuration Mode	Config Router BGP	
History	3.4.0000	First version
	3.6.3004	Added “disable” parameter
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 transport connection-mode passive	
Related Commands		
Note	• 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	

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 is 1-4094.
Default	BGP uses best local address	
Configuration Mode	Config Router BGP	
History	3.3.5006	First version
	3.3.5200	Updated example
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.2 update-source vlan 10	
Related Commands	N/A	
Note		

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 First version
Role	admin
Example	switch (config router bgp 100)# neighbor 10.10.10.2 no-update-source
Related Commands	N/A
Note	

neighbor weight

neighbor {<ip-address> | <peer-group-name>} weight <value>
no neighbor {<ip-address> | <peer-group-name>} weight

Assigns a weight attribute to paths from the specified neighbor.
 The no form of the command resets to default values.

Syntax Description	ip-address	IP address of the neighbor
	peer-group-name	Peer group name
	value	Weight value • Range: 1-65535 • Explicitly configuring a default value prevents the system from inheriting this parameter's group configuration
Default	Value is 32768 for router-originated paths and 0 for routes received through BGP	
Configuration Mode	Config Router BGP	
History	3.4.0000	First version
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 weight 100	
Related Commands	N/A	
Note	• 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.	

network

network <ip-prefix> <length> [<route-map-name>]
no network <ip-prefix> <length> [<route-map-name>]

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	A string that specific route map is assigned to the network.
	length	/24 or 255.255.255.0 format.
	route-map-name	The name of a route-map which is used to set the route’s attributes when it is advertised.
Default	N/A	
Configuration Mode	Config Router BGP	
History	3.3.5006	First version
	3.3.5200	Updated description, syntax description and notes
Role	admin	
Example	switch (config router bgp 100)# network 10.10.10.0 /24 routemap	
Related Commands		
Note	• 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) • This command cannot be used with route-maps	

redistribute

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
	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	
Role	admin	
Example	switch (config router bgp 100)# redistribute ospf	
Related Commands	N/A	
Note	- Multiple redistribution options can be applied - This command cannot be used with route-maps	

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
Role	admin	
Example	switch (config router bgp 100)# router-id 10.10.10.10	
Related Commands		
Note	The IP address configured identifies the BGP speaker. The command triggers an automatic notification and session reset for the BGP neighbors.	

6.3.4.3 Show

show ip bgp

show ip bgp [<ip-address> <mask> [detail | longer-prefixes [detail]]]

Displays information about the BGP routes table (RIB).

Syntax Description	ip-address	IP address (e.g. 172.3.12.4).
	mask	Netmask (e.g. /24 or 255.255.255.0).
	detail	Displays detailed information about a subset of the bgp learned routes.
	longer-prefixes	Displays the routes to the specified destination and any routes to a more specific destination. Example: If “10.20.30.0 /24 longer-prefixes” is run, all routes starting with 10.20.30 regardless of the prefix length (10.20.30.X /24, 10.20.30.X /25, etc.) are displayed – providing there are any such routes received/sent from/to that neighbor.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.5200	
Role	admin	
Example	<pre> switch (config) # show ip bgp BGP table version is 100, local router ID is 16.0.1.1 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 100.100.100.0/24 2.2.2.2 0 2 50 100 e 100.100.100.0/24 2.2.2.12 0 12 50 100 e Network Next Hop Metric LocPrf Weight Path 20.20.20.0/24 2.2.2.2 0 2 20 e 40.40.40.0/24 4.4.4.4 0 4 40 i 100.100.90.32/28 2.2.2.2 0 2 100 i 100.100.100.0/24 4.4.4.4 0 4 50 i switch (config) # </pre>	
Related Commands	N/A	
Note		

show ip bgp community

show ip bgp community <comm₁> <comm₂> ... <comm_n> [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
Role	admin
Example	<pre>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</pre>
Related Commands	N/A
Note	

show ip bgp neighbors

show ip bgp neighbors

Displays summaries information about all BGP neighbors.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.5200
Role	admin
Example	<pre>switch (config) # show ip bgp neighbors 3.5.7.5 received BGP table version is 66, 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 *> 100.0.20.0/24 3.5.7.5 10 100 0 5 i *> 3.5.7.128/32 3.5.7.5 7 100 0 5 i *> 100.0.30.0/24 3.5.7.5 0 100 0 5 i *> 10.20.30.0/24 3.5.7.5 0 100 0 5 12 i switch (config) #</pre>
Related Commands	N/A
Note	

show ip bgp neighbors <ip>

show ip bgp neighbors <ip-address>

Displays BGP summary information.

Syntax Description	ip-address	Neighbor IP address.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.5200	
Role	admin	
Example	<pre>switch (config) # show ip bgp neighbors 3.5.7.5 received BGP table version is 66, 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 *> 100.0.20.0/24 3.5.7.5 10 100 0 5 i *> 3.5.7.128/32 3.5.7.5 7 100 0 5 i *> 100.0.30.0/24 3.5.7.5 0 100 0 5 i *> 10.20.30.0/24 3.5.7.5 0 100 0 5 12 i switch (config) #</pre>	
Related Commands	N/A	
Note		

show ip bgp neighbors <ip> received

show ip bgp neighbors <ip-address> received [<ip-address> [<mask>] [longer-prefixes]]

Displays BGP summary information.

Syntax Description	ip-address	Neighbor IP address
	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.)
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.5200	
Role	admin	
Example		
Related Commands	N/A	
Note		

show ip bgp paths

show ip bgp paths

Displays summary of all AS paths.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.5200
Role	admin
Example	<pre>switch (config) # show ip bgp paths Refcount Metric Path 1 0 4 50 100 1 0 2 50 100 1 0 4 40 1 0 12 50 100 1 0 2 1 0 2 20 switch (config) #</pre>
Related Commands	N/A
Note	

show ip bgp peer-group

show ip bgp peer-group [<peer-group-name>]

Displays information about peer groups.

Syntax Description	peer-group-name	Displays information about a specific peer-group.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.4.0000	
Role	admin	
Example	<pre> switch (config) # show ip bgp peer-group BGP Peer-group [grpA]: Hold time: 1, Keep-alive: 60 Allow as-in: 0 Weight: 32768 Max prefix: 12000 Export local preferences: 100, Import local preferences: 100 Soft reconfiguration: set Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd 3.5.7.5 0 5 0 0 0 0 0 0:00:00:42 CONNECT 100.100.100.100 0 100 0 0 0 0 0 Never IDLE BGP Peer-group [grpB]: Hold time: 1, Keep-alive: 60 Allow as-in: 0 Weight: 32768 Max prefix: 12000 Export local preferences: 100, Import local preferences: 100 Soft reconfiguration: set Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd 3.4.3.7 0 7 0 0 0 0 0 0:00:00:17 ACTIVE BGP Peer-group [tomor_group]: Hold time: 1, Keep-alive: 60 Allow as-in: 0 Weight: 32768 Max prefix: 12000 Export local preferences: 100, Import local preferences: 100 Soft reconfiguration: set Peer-groups count: 3 switch-e07c04 [standalone: master] (config) # </pre>	

Related Commands	N/A

Note

show ip bgp summary

show ip bgp summary

Displays BGP summary information.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.5200
Role	admin
Example	<pre>switch (config) # show ip bgp summary BGP router identifier 3.5.7.4, local AS number 4 BGP table version is 70, main routing table version 70 8 network entries using 2176 bytes of memory 4 path entries using 1088 bytes of memory 4 BGP path attribute entries using 256 bytes of memory 0 multipath network entries and 0 multipath paths 4 BGP community entries using 64 bytes of memory 0 received paths for inbound soft reconfiguration BGP using 26308 total bytes of memory Dampening disabled. 0 history paths, 0 dampened paths BGP activity 37/8 prefixes, 37/4 paths Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd 3.4.3.7 4 7 3 9 70 0 0 0:00:00:48 ESTABLISHED 3.5.7.5 0 5 0 0 0 0 0 0:00:01:54 CONNECT 100.100.100.100 0 100 0 0 0 0 0 Never IDLE switch-e07c04 [standalone: master] (config) #</pre>
Related Commands	N/A
Note	

6.3.5 IP AS-Path Access-List

6.3.5.1 Commands

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	Regular expression that is used to specify a pattern to match against an input string.
	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	
Role	admin	
Example	switch (config)# ip as-path access-list mylist permit	
	switch (config)#	
Related Commands	N/A	
Note	If access list_name does not exist, this command creates it. If it already exists, this command appends statements to the list.	

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	
Role	admin	
Example	switch (config)# show ip as-path access-list mylist	
Related Commands	N/A	
Note		

6.3.6 IP Community-List

6.3.6.1 Commands

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	
Role	admin	
Example	switch (config)# ip community-list standard mycommunity permit 1:2 3:4	
Related Commands	N/A	
Note	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.	

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	
Role	admin	
Example	switch (config)# ip community-list expanded mycommunity permit 1:[0-9]+	
Related Commands	N/A	
Note	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.	

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	
Role	admin	
Example	switch (config)# show ip community-list mycommunity	
Related Commands	N/A	
Note	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.	

6.4 Policy Rules

6.4.1 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](#)” on [page 1159](#) or “[redistribute](#)” on [page 1160](#).

6.4.1.1 Commands

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.

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	
	3.3.5200	Updated notes
Role	admin	
Example	switch (config) # route-map mymap permit 1200 switch (config route-map mymap permit 1200)#	
Related Commands	N/A	
Note	- 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.	

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	sequence-number	
Default	N/A	
Configuration Mode	Config Route Map	
History	3.3.5006	First version
	3.3.5200	Updated example
Role	admin	
Example	<pre> 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 switch (config)# show route-map test route-map test, permit, sequence 10 Match clauses: as-number 40 Set clauses: weight 7 continue 1200 switch (config route-map mymap permit 10)# route-map test permit 10 no continue switch (config route-map mymap permit 10)# show route-map test route-map test, permit, sequence 10 Match clauses: as-number 40 Set clauses: weight 7 switch (config route-map mymap permit 10)# </pre>	

Related Commands`route-map <map-name> [deny | permit] [sequence-number]`**Note**

- 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.
-
-

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 First version 3.3.5200 Updated example
Role	admin
Example	<pre> switch (config)# route-map mymap permit 10 match as-number 40 switch (config)# route-map mymap permit 10 set weight 7 switch (config)# show route-map test route-map test, permit, sequence 10 Match clauses: as-number 40 Set clauses: weight 7 switch (config)# route-map mymap permit 1200 switch (config route-map mymap permit 1200)# set weight 11 switch (config route-map mymap permit 1200)# abort switch (config)# show route-map mymap route-map mymap, permit, sequence 10 Match clauses: as-number 40 Set clauses: weight 7 switch (config)# </pre>
Related Commands	N/A
Note	

exit

exit

Saves pending route map clause changes to running-config and returns to global configuration mode.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config Route Map
History	3.3.5006
Role	admin
Example	<pre> switch (config)# route-map mymap permit 10 match as-number 40 switch (config)# route-map mymap permit 10 set weight 7 switch (config)# show route-map test route-map test, permit, sequence 10 Match clauses: as-number 40 Set clauses: weight 7 switch (config)# route-map mymap permit 1200 switch (config route-map mymap permit 1200)# set weight 11 switch (config route-map mymap permit 1200)# exit switch (config)# show route-map test 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 switch (config)# </pre>
Related Commands	N/A
Note	

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	
Role	admin	
Example	<pre>switch (config route-map mymap permit 10)# match as-number 40 switch (config route-map mymap permit 10)#</pre>	
Related Commands	N/A	
Note	- 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.	

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
Role	admin	
Example	switch (config route-map mymap permit 10)# match as-path my-list	
Related Commands	N/A	
Note	- 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	

match community

match community <list-of-communities> [exact-match]
no match community <list-of-communities>

Creates a route map clause entry that matches a route if it contains at least the specified communities.
 The no form of the command removes the match clause.

Syntax Description	list of communities	List of standard communities: • <aa:nn> • <number> • internet • local-AS • no-advertise • no-export
	exact-match	Creates a route map clause entry that matches the route's communities exactly.
Default	N/A	
Configuration Mode	Config Route Map	
History	3.3.5006	
Role	admin	
Example	switch (config route-map mymap permit 10)# match community 1:100 3:52	
Related Commands	N/A	
Note	• 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. • Route-map's match on a list of communities is performed with the command "match community-list" and not this command.	

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
Role	admin
Example	switch (config route-map mymap permit 10)# match community-list COM_LIST exact-match
Related Commands	N/A
Note	- 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.

match interface

match interface <interface-type> <number>
no match interface

Matches the route's interface
 The no form of the command removes the match clause.

Syntax Description	prefix-list-name	Prefix-list name.
Default	N/A	
Configuration Mode	Config Route Map	
History	3.3.5006	
Role	admin	
Example	switch (config route-map mymap permit 10)# match interface ethernet 1/1	
Related Commands	N/A	
Note	- 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.	

match ip address

match ip address <prefix-list-name>
no match ip address

Filters according to IPv4 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	
Role	admin	
Example	switch (config route-map mymap permit 10)# match ip address listSmallRoutes	
Related Commands	N/A	
Note	- 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.	

match ip next-hop

match ip next-hop <value>
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	value	Next hop IP address: A.B.C.D (e.g. 10.0.13.86).
Default	N/A	
Configuration Mode	Config Route Map	
History	3.3.5200	
Role	admin	
Example	switch (config route-map mymap permit 10)# match ip next-hop 10.10.10.10	
Related Commands	N/A	
Note	- 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.	

match local-preference

match local-preference <value>
no match local-preference

Configures a route's entry local-preference match.
The no form of the command removes a route-map's entry local-preference match.

Syntax Description	value	Range: 1-2147483647.
Default	N/A	
Configuration Mode	Config Route Map	
History	3.3.5200	First version
	3.4.0000	Updated value range
Role	admin	
Example	switch (config route-map mymap permit 10)# match local-preference 10	
Related Commands	N/A	
Note	- 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.	

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	First version
	3.4.0000	Updated value range
Role	admin	
Example	switch (config route-map mymap permit 10)# match metric 10	
Related Commands	N/A	
Note	- 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.	

set as-path prepend

set as-path prepend <value₁> <value₂> ... <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	
Role	admin	
Example	switch (config route-map mymap permit 10)# set as-path prepend 5 10	
Related Commands	N/A	
Note		

set as-path tag

set as-path tag <value>
no set as-path tag

Configures a route's entry AS-path tag parameter.
The no form of the command removes a route-map's entry AS path tag setting.

Syntax Description	value	Range: 1-2147483648.
Default	N/A	
Configuration Mode	Config Route Map	
History	3.3.5200	
Role	admin	
Example	switch (config route-map mymap permit 10)# set as-path tag 1	
Related Commands	N/A	
Note		

set community

set community {<list of communities> | none}
no set community {<list of communities> | none}

Sets the community attribute of a distributed route
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	
Role	admin	
Example	switch (config route-map mymap permit 10)# set community 1:2 3:4	
Related Commands	N/A	
Note		

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	
Role	admin	
Example	switch (config route-map mymap permit 10)# set community none	
Related Commands	N/A	
Note		

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
Role	admin
Example	switch (config route-map mymap permit 10)# set community none
Related Commands	N/A
Note	

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	
Role	admin	
Example	<pre>switch-e07c04 [standalone: master] (config) # route-map test_route_map switch-e07c04 [standalone: master] (config route-map test_route_map permit 10) # set community 400:1 delete</pre>	
Related Commands	N/A	
Note		

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	
	3.6.3004	Added note
Role	admin	
Example	switch (config route-map mymap permit 10)# set community internet 1:3 additive	
Related Commands	N/A	
Note	A community-list must already exist before a node is configured to use it	

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
Role	admin
Example	switch (config route-map mymap permit 10)# set community-list mycommunity additive
Related Commands	N/A
Note	

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	
Role	admin	
Example	switch (config route-map mymap permit 10)# set community-list mycommunity delete	
Related Commands	N/A	
Note		

set ip next-hop

set ip next-hop <value>
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	value	Route next-hop IP: A.B.C.D (e.g. 10.0.13.86).
Default	N/A	
Configuration Mode	Config Route Map	
History	3.3.5200	
Role	admin	
Example	switch (config route-map mymap permit 10)# set ip next-hop 10.10.10.10	
Related Commands	N/A	
Note		

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: 1-2147483648.
Default	N/A	
Configuration Mode	Config Route Map	
History	3.3.5200	
Role	admin	
Example	switch (config route-map mymap permit 10)# set local-preference 10	
Related Commands	N/A	
Note		

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: 1-2147483647.
Default	N/A	
Configuration Mode	Config Route Map	
History	3.3.5200	
Role	admin	
Example	switch (config route-map mymap permit 10)# set metric 10	
Related Commands	N/A	
Note		

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	
Role	admin	
Example	switch (config route-map mymap permit 10)# set origin egp	
Related Commands	N/A	
Note		

set tag

set tag <value>
no set tag

Configures a route's entry tag parameter.
 The no form of the command removes a route-map's entry tag setting.

Syntax Description	value	Range: 1-2147483647.
Default	N/A	
Configuration Mode	Config Route Map	
History	3.3.5200	
	3.4.0000	Updated parameter range
Role	admin	
Example	switch (config route-map mymap permit 10)# set tag 10	
Related Commands	N/A	
Note		

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	First version
	3.4.0000	Updated parameter range
Role	admin	
Example	switch (config route-map mymap permit 10)# set weight 7	
Related Commands	route-map <map-name> [deny permit] [sequence-number]	
Note		

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
Role	admin
Example	<pre>switch (config)# show route-map mymap route-map mymap, permit, sequence 1200 Set clauses: continue 1800 switch (config)#</pre>
Related Commands	N/A
Note	

6.4.2 IP Prefix-List

Prefix-list is a list of entries, each of which can match one or more IP prefixes. A prefix-list is usually used to match a specific IP prefix, mostly in relation to IP route destinations.

The prefix is considered to match the list if one of the entries match the prefix; the entry itself can be marked as a “permit” entry or a “deny” entry, which can be used by the matching code to decide if the route is to be accepted or not.

The prefix is matched to the prefix-list entries in the order of the sequence number of the entries in the list.

6.4.2.1 Commands

ip prefix-list

```
ip prefix-list <list-name> [seq <number>] {permit | deny} <ip> [eq <length> |
<prefix> [eq <length> | le <length> | ge <length> [le <length>]]]
no ip prefix-list <list-name> [seq <number>]
```

Creates or updates a prefix-list.

The no form of the command deletes a prefix-list or a prefix-list entry

Syntax Description	list-name	String
	seq <number>	Sequence number assigned to entry. Range: 0-65535.
	permit	Permits access for a matching condition.
	deny	Denies access for a matching condition.
	ip	IP 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	Sequence value = 10	
Configuration Mode	Config	
History	3.3.5200	
Role	admin	
Example	<pre>switch (config)# ip prefix-list a-list permit 10.20.0.0 /16 eq 24 switch (config)#</pre>	
Related Commands	N/A	
Note		

show ip prefix-list

show ip prefix-list [<name>]

Displays prefix-lists.

Syntax Description	name	Displays a specific prefix-list.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.5200	
Role	admin	
Example	<pre>switch (config)# show ip prefix-list prefix-list: a-list count: 1, range entries: 1, sequences: 10 - 10 seq 10 permit 10.20.0.0 /16 ge 24 (hit count: 0, refcount: 0) prefix-list: b-list count: 2, range entries: 2, sequences: 10 - 20 seq 10 deny 10.10.0.0 /16 le 24 (hit count: 0, refcount: 0) seq 20 deny 10.20.0.0 /16 le 24 (hit count: 0, refcount: 0) switch (config)#</pre>	
Related Commands	N/A	
Note		

6.5 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 PIM dense mode (PIM-DM), PIM sparse mode (PIM-SM, which is not supported on Mellanox platforms), Bidirectional PIM (PIM-BIDIR) and Bootstrap router (BSR) protocol.

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.

6.5.1 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.

6.5.2 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.

6.5.2.1 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.

6.5.2.2 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.

6.5.3 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.

6.5.4 Configuring Multicast

Precondition steps:

Step 1. Enable IP routing functionality. Run:

```
switch (config)# ip routing
```

Step 2. Enable the desired VLAN. Run:

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

Step 3. Add this VLAN to the desired interface. Run:

```
switch (config)# interface ethernet 1/1  
switch (config ethernet 1/1)#switchport access vlan 10
```

Step 4. Create a VLAN interface. Run:

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

Step 5. Apply IP address to the VLAN interface. Run:

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

Step 6. Enable the interface. Run:

```
switch (config interface vlan 10)# no shutdown
```

6.5.4.1 Configuring IGMP

IGMP is enabled when IP multicast is enabled and static multicast or PIM is enabled on the interface.

6.5.4.2 Verifying IGMP

Step 1. Display a brief IGMP interface status. Run:

```
switch (config)# show ip igmp interface brief  
IGMP Interfaces for VRF "default", Count: 1  
Interface      IP Address      IGMP Querier    Membership      Version  
VLAN10         10.10.10.1      10.10.10.1      5               v2
```

Step 2. Display detailed IGMP interface status. Run:

```
switch (config)#show ip igmp interface vlan 10
IGMP Interfaces for VRF "default"

VLAN10
Interface status: protocol-up/admin-up/link-up
IP address: 10.10.10.1, IP Subnet: 10.10.10.0/24
Active Querier: 10.10.10.1
Membership count: 5
Route-queue depth: 0
IGMP Version: 2
IGMP query interval: 125 secs, configured value: 125 secs
IGMP max response time: 10 secs, configured value: 10 secs
IGMP startup query interval: 125 secs, configured value: 125 secs
IGMP startup query count: 2
IGMP group timeout: 260 secs, configured value: 260 secs
IGMP querier timeout: 260 secs configured value: 260 secs
IGMP last member mrt: 25 secs configured value: 25
IGMP robustness variable: 2
IGMP interface immediate leave: Disabled
IGMP interface statistics:
General (sent/received):
v1/v2-reports: 0/10
v2-queries: 271/0, v2-leaves: 0/0
v3-queries: 0/0,
v3-reports: 0/0
switch (config)#
```

Step 3. Display the list of IGMP groups and their status. Run:

```
switch (config)#show ip igmp groups
IGMP Connected Group Membership for VRF "default", - 2 total entries
Type: S - Static, D - Dynamic, L - Local, T - SSM Translated
Group Address Type Interface Uptime Expires Last
Reporter
226.0.1.0 D vlan10 [0d 00:00:07.46] [0d 00:04:05.08] 10.10.10.2
226.0.1.1 D vlan10 [0d 00:00:07.47] [0d 00:04:05.08]
10.10.10.2
switch (config)#
```

6.5.4.3 Configuring PIM

Prerequisites:

Step 1. If not enabled, enable IP routing. Run:

```
switch (config)# ip routing
```

Step 2. Globally enable multicast routing. Run:

```
switch (config)# ip multicast-routing
```

➤ **To configure PIM:**

Step 1. Enable PIM. Run:

```
switch (config)# protocol pim
```

Step 2. Globally enable Bidirectional PIM (BIDIR mode). Run:

```
switch (config)# no ip pim bidir shutdown
```

6.5.5 Commands

6.5.5.1 PIM

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
Role	admin
Example	switch (config) # protocol pim
Related Commands	N/A
Note	

ip pim bidir shutdown

ip pim bidir shutdown
no ip pim bidir shutdown

Disables PIM bidir.
The no form of the command enables PIM bidir.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config
History	3.3.5006
Role	admin
Example	switch (config) # no ip pim bidir shutdown
Related Commands	N/A
Note	

ip pim rp-address

ip pim rp-address <rp-address> [group-list <ip-address> <prefix>] [override] bidir
no ip pim rp-address <rp-address> [group-list <ip-address> <prefix>]

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	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	Specifies that the group range uses a bidirectional PIM.
Default	N/A	
Configuration Mode	Config	
History	3.3.5006	
Role	admin	
Example	switch (config) # ip pim rp-address 10.10.10.10 bidir	
Related Commands	N/A	
Note		

ip pim bsr-candidate

```
ip pim bsr-candidate {vlan <vlan-id> | loopback <number> | ethernet <port>}
[hash-len <hash-length>] [priority <priority>] [interval <interval>]
no ip pim 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	vlan <vlan-id>	The VLAN ID. Range is 1-4094.
	loopback <number>	Loopback interface number.
	ethernet <port>	Ethernet interface.
	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	
Role	admin	
Example	switch (config) # ip pim bsr-candidate vlan 10 priority 100	

Related Commands `ip pim sparse-mode`**Note**

- IP PIM sparse-mode must be enabled on the interface.
 - 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.
-

ip pim bsr-holdtime

ip pim bsr-holdtime <period>
no ip pim bsr-holdtime

Configures the timeout period an elected BSR remains valid after receiving a BSM.

The no form of the command resets the parameters to their default.

Syntax Description	period	In seconds. Range: 12-1073741823 (1.073 billion).
Default	period = 2*(BSR candidate interval) + 10	
Configuration Mode	Config	
History	3.3.5006	
Role	admin	
Example	switch (config) # ip pim bsr-holdtime 30	
Related Commands		
Note		

ip pim rp-candidate

```
ip pim rp-candidate {vlan <vlan-id> | loopback <number> | ethernet <slot/
port>} group-list <ip-address> <prefix> [bidir] [priority <priority>] [interval
<interval>]
no ip pim rp-candidate {vlan <vlan-id> | loopback <number> | ethernet <slot/
port>} group-list <ip-address> <prefix> [bidir] [priority <priority>] [interval
<interval>]
```

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	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.
Default	The RP priority is 192. The BSR message interval is 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	
Role	admin	
Example	switch (config) # ip pim rp-candidate vlan 19 group-list 225.6.5.0 /25 priority 20 interval 30 bidir	

Related Commands

N/A

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. The default interval value is 60 seconds.
 - 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.
-

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
Role	admin
Example	switch (config interface vlan 10) # ip pim sparse-mode
Related Commands	N/A
Note	

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	
Role	admin	
Example	switch (config interface vlan 10) # ip pim dr-priority 5	
Related Commands	ip pim sparse-mode	
Note	The command “ip pim sparse-mode” must be run prior to using this command.	

ip pim hello-interval

ip pim hello-interval <interval>
no ip pim hello-interval

Configures PIM Hello interval in milliseconds.
The no form of the command resets this parameter to its default.

Syntax Description	interval	PIM Hello interval in milliseconds. Range:1000-65535000.
Default	30,000 milliseconds	
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	
Role	admin	
Example	switch (config interface vlan 10) # ip pim hello-interval 70000	
Related Commands	ip pim sparse-mode	
Note	The command “ip pim sparse-mode” must be run prior to using this command.	

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: 1-1000000 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	
Role	admin	
Example	switch (config interface vlan 10) # ip pim join-prune-interval 60	
Related Commands		
Note		

ip pim border

ip pim border
no ip pim border

Configures an interface on an IPv4 PIM border.
The no form of the command removes the interface from being a PIM border.

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
Role	admin
Example	switch (config interface vlan 10) # ip pim border
Related Commands	
Note	PIM border blocks PIM control traffic, but sends and receives all multicast traffic.

ip pim bsr-border

ip pim bsr-border
no ip pim bsr-border

Prevents the switch from sending bootstrap router messages (BSMs) over the configuration mode interface.
 The no form of the command resets the parameter to its default value.

Syntax Description	N/A
Default	no pim bsr-border
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
Role	admin
Example	switch (config interface vlan 10) # ip pim bsr-border
Related Commands	
Note	

ip pim multipath rp

ip pim multipath rp
no ip pim multipath rp

Enables PIM load-sharing for Rendezvous Points (RPs).
The no form of the command disables PIM load-sharing for RPs.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config
History	3.4.2008
Role	admin
Example	switch (config) # ip pim multipath rp
Related Commands	N/A
Note	

debug ethernet ip pim

```
debug ethernet ip pim {all | control-plane | data-path | fail-all | init-shut |
management | memory | packet-dump | resources}
no debug ethernet ip pim {all | control-plane | data-path | fail-all | init-shut |
management | memory | packet-dump | resources}
```

Configures the trace level for PIM.

The no form of the command removes the trace level for PIM.

Syntax Description	all	Enable track traces.
	control-plane	Control plane traces.
	data-path	IP packet dump trace.
	fail-all	All failures including Packet Validation Trace.
	init-shut	Init and shutdown messages.
	management	Management messages.
	memory	Memory related messages.
	packet-dump	Packet dump messages.
	resources	OS Resource trace.
Default	N/A	
Configuration Mode	Config	
History	3.3.5200	
Role	admin	
Example	switch (config)# debug ethernet ip pim all	
Related Commands		
Note		

show ip pim protocol

show ip pim protocol

Displays PIM protocol information (counters).

Syntax Description	N/A																																															
Default	N/A																																															
Configuration Mode	Any Command Mode																																															
History	3.3.5200																																															
Role	admin																																															
Example	<pre>switch (config) # show ip pim protocol PIM Control Counters</pre> <table><thead><tr><th></th><th>Received</th><th>Sent</th><th>Invalid</th></tr></thead><tbody><tr><td>Assert</td><td>0</td><td>0</td><td>0</td></tr><tr><td>Bootstrap Router</td><td>0</td><td>0</td><td>0</td></tr><tr><td>CRP Advertisement</td><td>0</td><td>0</td><td>0</td></tr><tr><td>Graft</td><td>0</td><td>0</td><td>0</td></tr><tr><td>Grapt Ack</td><td>0</td><td>0</td><td>0</td></tr><tr><td>Hello</td><td>0</td><td>0</td><td>0</td></tr><tr><td>J/P</td><td>0</td><td>0</td><td>0</td></tr><tr><td>Register</td><td>0</td><td>0</td><td>0</td></tr><tr><td>Register Stop</td><td>0</td><td>0</td><td>0</td></tr><tr><td>State Refresh</td><td>0</td><td>0</td><td>0</td></tr></tbody></table> <pre>switch (config) #</pre>					Received	Sent	Invalid	Assert	0	0	0	Bootstrap Router	0	0	0	CRP Advertisement	0	0	0	Graft	0	0	0	Grapt Ack	0	0	0	Hello	0	0	0	J/P	0	0	0	Register	0	0	0	Register Stop	0	0	0	State Refresh	0	0	0
	Received	Sent	Invalid																																													
Assert	0	0	0																																													
Bootstrap Router	0	0	0																																													
CRP Advertisement	0	0	0																																													
Graft	0	0	0																																													
Grapt Ack	0	0	0																																													
Hello	0	0	0																																													
J/P	0	0	0																																													
Register	0	0	0																																													
Register Stop	0	0	0																																													
State Refresh	0	0	0																																													
Related Commands																																																
Note																																																

show ip pim bsr

show ip pim bsr

Displays PIM BSR information.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.5006
Role	admin
Example	<pre>arc-switch14 [standalone: master] (config) # show ip pim bsr PIMv2 Bootstrap information BSR address: 4.4.4.14 Uptime: 00:00:30, BSR Priority: 0, Hash mask length: 30 Expires: 00:00:57 This system is a candidate BSR Candidate BSR address: 4.4.4.14, priority: 0, hash mask length: 30 interval: 60, holdtime: 130</pre>
Related Commands	
Note	

show ip pim neighbor

show ip pim neighbor [vlan <vlan-id> | <other interfaces> | <ip-addr>]

Displays information about IPv4 PIM neighbors.

Syntax Description	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	
Role	admin	
Example	<pre>switch (config) # show ip pim neighbor PIM Neighbor Status for VRF "default" Neighbor Interface Uptime Expires Ver DR Prio Mode 5.5.5.1 VLAN5 10:36:45 00:01:43 1 9.9.9.1 VLAN9 10:36:42 00:01:43 1 switch (config) #</pre>	
Related Commands		
Note		

show ip pim rp

show ip pim rp <rp-address>

Displays information about the rendezvous points (RPs) for PIM.

Syntax Description	rp-address	A rendezvous points address.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.5006	
Role	admin	
Example	<pre>switch(config)# show ip pim rp PIM RP Status Information for VRF "default" BSR: 10.10.10.10, expires: 00:01:16, priority: 255, hash-length: 0 RP: 11.11.11.11, expires: 00:01:36 priority: 0, RP-source: 10.10.10.10, group ranges: 225.10.0.0/24 RP: 8.8.8.2, expires: 00:01:36 priority: 0, RP-source: 10.10.10.10, group ranges: 225.12.0.0/24 switch(config)#</pre>	
Related Commands		
Note		

show ip pim rp-hash

show ip pim rp-hash <group>

Displays the hashed value of the group (RP address according the group address).

Syntax Description	group	Filters the output per a specific IP Multicast group address.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.5006	
Role	admin	
Example	<pre>switch (config) # show ip pim rp-hash 225.7.6.2 RP 20.20.20.49, v2 Info Source: 20.20.20.49, via bootstrap, priority 60, holdtime 57 Expires: 00:00:53 PIMv2 Hash Value (mask 255.255.255.252) switch (config)#</pre>	
Related Commands		
Note		

show ip pim rp-candidate

show ip pim rp-candidate

Displays information about RP candidate status.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.5006
Role	admin
Example	<pre>switch (config)# show ip pim rp-candidate Next Candidate-RP-Advertisement in 00:11:22/00:60:00 RP: 10.10.10.10 group prefixes priority 224.0.0.0/4 190 225.0.0.0/4 191 switch (config)#</pre>
Related Commands	
Note	

show ip pim interface

show ip pim interface {[vlan <vlan id> | ethernet <port>] [df] | brief}

Displays information about the enabled interfaces for PIM.

Syntax Description	vlan <vlan-id>	Filters the output for specific interface.
	ethernet <port>	Ethernet interface.
	df	Displays information about elected designated forwarders.
	brief	Displays a summary of information for all interfaces.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.5006	
Role	admin	
Example	<pre># arc-switch55 [standalone: master] (config) # show ip pim interface vlan 2919 Interface Vlan2919 address is 70.28.23.80 PIM: enabled PIM version: 2, mode: sparse PIM DR: 70.28.23.80 (this system) PIM DR Priority: 1 PIM configured DR priority: PIM neighbor count: 1 PIM neighbor holdtime: 105 secs PIM Hello Interval: 30 seconds, next hello sent in: 00:00:28 PIM Hello Generation ID: 61345 PIM Join-Prune Interval: 60 seconds PIM domain border: no PIM Interface Statistics: General (sent/received): Hellos: 36/37, JPs: 0/0, Asserts: 0/0 Grafts: 0/0, Graft-Acks: 0/0 DF-Offers: 0/0, DF-Winners: 0/0, DF-Backoffs: 0/0, DF- Passes: 0/0 Errors: Checksum errors: 0, Invalid packet types/DF subtypes: 0/0 Authentication failed: 0 Packets from non-neighbors: 1 JPs received on RPF-interface: 0 (*,G) Joins received with no/wrong RP: 0/0 (*,G)/(S,G) JPs received for SSM/Bidir groups: 0/0</pre>	

Related Commands

Note

show ip pim upstream joins

show ip pim upstream joins

Displays information about any PIM joins/prunes which are currently being sent to upstream PIM routers

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.5006
Role	admin
Example	<pre>switch (config) # show ip pim upstream joins Neighbor address: 159.135.45.26 via interface: 159.135.45.34 next message in 43 seconds Group: 224.0.10.0 Joins: 22.74.49.25 Prunes: switch (config) #</pre>
Related Commands	
Note	Should contain the following information: neighbor address, interface address, group range, Joins, Prunes.

6.5.5.2 Multicast

ip multicast-routing

ip multicast-routing
no ip multicast-routing

Allows the switch to forward multicast packets.
The no form of the command disables multicast routing.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config
History	3.3.5006
Role	admin
Example	switch (config)# ip multicast-routing
Related Commands	N/A
Note	

ip mroute

ip mroute {<ip-addr> <ip-mask> <next-hop>} [pref]
no ip mroute {<ip-addr> <ip-mask>}

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	
Role	admin	
Example	arc-switch14 [standalone: master] (config) # ip mroute 16.16.0.0 /16 3.3.3.1	
Related Commands	N/A	
Note		

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	
Role	admin	
Example	switch (config interface vlan 10)# ip multicast ttl-threshold 10	
Related Commands	N/A	
Note		

show ip mroute

show ip mroute [summary | <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.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.2.1000	
	3.5.1000	Added new F flag and updated Example
Role	admin	
Example	<pre>switch (config) # show ip mroute IP Multicast Routing Table Flags: B - Bidir Group, L - Local, P - Pruned, R - RP-bit set, T - SPT-bit set J - Join SPT, F - Failed to install in H/W Timers: Uptime/Expires Interface state: Interface, State/Mode (*, 234.10.0.0/16), 00D 01:06:04, RP 10.10.10.10, flags: BR Bidir-Upstream: Eth1/10 Outgoing interface list: Eth1/10, Forwarding/Sparse, 00D 01:06:04/00D 00:00:00 F(*, 234.8.0.0/16), 00D 01:06:03, RP 10.10.10.10, flags: BR Bidir-Upstream: Eth1/10 Outgoing interface list: Eth1/10, Forwarding/Sparse, 00D 01:06:04/00D 00:00:00</pre>	
Related Commands	N/A	
Note		

6.5.5.3 IGMP

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	3.3.5006
Role	admin
Example	switch (config interface vlan 10)# ip igmp immediate-leave
Related Commands	N/A
Note	

ip igmp last-member-query-count

ip igmp last-member-query-count <count>
no ip igmp last-member-query-count

Configures the number of query messages the switch sends in response to a group-specific or group-source-specific leave message.
The no form of the command resets this parameter to its default.

Syntax Description	Count	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	
Role	admin	
Example	switch (config interface vlan 10)# ip igmp last-member-query-count 7	
Related Commands	N/A	
Note	This parameter reflects expected packet loss on a congested network.	

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	
Role	admin	
Example	switch (config interface vlan 10)# ip igmp last-member-query-response-time 10	
Related Commands	N/A	
Note		

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: 1-65535.
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	
Role	admin	
Example	switch (config interface vlan 10)# ip igmp startup-query-count 10	
Related Commands	N/A	
Note		

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	30	
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	
Role	admin	
Example	switch (config interface vlan 10)# ip igmp startup-query-interval 10	
Related Commands	N/A	
Note		

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	
Configuration Mode	Config Interface VLAN	
History	3.3.5006	
Role	admin	
Example	switch (config interface vlan 10)# ip igmp query-interval 60	
Related Commands	N/A	
Note		

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	
Role	admin	
Example	switch (config interface vlan 10)# ip igmp query-max-response-time 20	
Related Commands	N/A	
Note		

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	
Role	admin	
Example	switch (config interface vlan 10)# ip igmp robustness-variable 4	
Related Commands	N/A	
Note	- 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.	

ip igmp static-oif

ip igmp static-oif <group>
no ip igmp static-oif

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.
Default	no ip igmp static-oif	
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	
Role	admin	
Example	switch (config interface vlan 10)# ip igmp static-oif 10.10.10.5	
Related Commands	N/A	
Note	PIM must be enabled in order to configure the route in the hardware.	

clear ip igmp groups

clear ip igmp groups {all | <group-address> <mask>}

Clears IGMP group information.

Syntax Description	all	Clears all IGMP groups.
	group-address	Clears a specific group.
Default	no ip igmp static-oif	
Configuration Mode	Config	
History	3.3.5200	
Role	admin	
Example	switch (config)# clear ip igmp groups all switch (config)#	
Related Commands	N/A	
Note		

debug ethernet ip igmp-l3

debug ethernet ip igmp-l3 {all | control-plane | data-path | fail-all | init-shut | management | memory | packet-dump | resources}
no debug ethernet ip igmp-l3 {all | control-plane | data-path | fail-all | init-shut | management | memory | packet-dump | resources}

Configures the trace level for IGMP.
 The no form of the command removes the trace level for IGMP.

Syntax Description	all	Enable track traces.
	control-plane	Control plane traces.
	data-path	IP packet dump trace.
	fail-all	All failures including Packet Validation Trace.
	init-shut	Init and shutdown messages.
	management	Management messages.
	memory	Memory related messages.
	packet-dump	Packet dump messages.
	resources	OS Resource trace.
Default	N/A	
Configuration Mode	Config	
History	3.3.5200	
Role	admin	
Example	switch (config)# debug ethernet ip igmp-l3 all	
Related Commands		
Note		

show ip igmp groups

show ip igmp groups [<group>] [vlan <vlan-id>]

Displays information about IGMP-attached group membership.

Syntax Description	group	Filters the output to a specific IP multicast group address.
	vlan <vlan-id>	Filters the output to a specific VLAN ID.
Default	N/A	
Configuration Mode	Any Command Mode	
History		
Role	admin	
Example	<pre>switch (config)# show ip igmp groups IGMP Connected Group Membership for VRF "default" Type: S - Static, D - Dynamic, L - Local, T - SSM Translated Group Address Type Interface Uptime Expires Last Reporter 225.7.6.0 S vlan19 [0d 00:12:12.14] [0d 00:00:00.00] 0.0.0.0 225.7.10.1 D vlan19 [0d 00:00:01.18] [0d 00:04:08.81] 19.19.19.1 225.7.7.7 S vlan19 [0d 00:12:12.15] [0d 00:00:00.00] 0.0.0.0 225.7.7.7 S vlan21 [0d 00:12:12.15] [0d 00:00:00.00] 0.0.0.0</pre>	
Related Commands	N/A	
Note		

show ip igmp interface

show ip igmp interface [vlan <vlan-id> | brief]

Displays IGMP brief configuration and status.

Syntax Description	brief	Displays brief output information.
	vlan <vlan-id>	Filters the output to a specific VLAN ID.
Default	N/A	
Configuration Mode	Any Command Mode	
History		
Role	admin	

Example

```
switch(config)#show ip igmp interface
IGMP Interfaces for VRF "default"

VLAN5
Interface status: protocol-down/admin-up/link-down
IP address: 5.5.5.49, IP Subnet: 5.5.5.0/24
Active Querier: 5.5.5.48
Membership count: 0
Route-queue depth: 0
IGMP Version: 2
IGMP query interval: 125 secs, configured value: 125 secs
IGMP max response time: 100 secs, configured value: 100 secs
IGMP startup query interval: 125 secs, configured value: 125 secs
IGMP startup query count: 2
IGMP group timeout: 350 secs, configured value: 350 secs
IGMP querier timeout: 350 secs configured value: 350 secs
IGMP last member mrt: 10 secs configured value: 10
IGMP robustness variable: 2
IGMP interface immediate leave: Disabled
IGMP interface statistics:
General (sent/received):
v1/v2-reports: 0/0
v2-queries: 3/1,v2-leaves: 0/0
v3-queries: 0/0,
v3-reports: 0/0

VLAN19
Interface status: protocol-up/admin-up/link-up
IP address: 19.19.19.49, IP Subnet: 19.19.19.0/24
Active Querier: 19.19.19.49
Membership count: 3
Route-queue depth: 0
IGMP Version: 2
IGMP query interval: 125 secs, configured value: 125 secs
IGMP max response time: 10 secs, configured value: 10 secs
IGMP startup query interval: 125 secs, configured value: 125 secs
IGMP startup query count: 2
IGMP group timeout: 260 secs, configured value: 260 secs
IGMP querier timeout: 260 secs configured value: 260 secs
IGMP last member mrt: 1 secs configured value: 1
IGMP robustness variable: 2
IGMP interface immediate leave: Disabled
IGMP interface statistics:
General (sent/received):
v1/v2-reports: 0/5
v2-queries: 14/0,v2-leaves: 0/1
v3-queries: 0/0,
v3-reports: 0/0
```

Related Commands

N/A

Note

6.6 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 sub-network.

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.

6.6.1 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.

Figure 41: Common VRRP Configuration with Load Balancing



6.6.2 Configuring VRRP

➤ To configure VRRP:

Precondition steps:

Step 1. Enable IP routing functionality. Run:

```
switch (config)# ip routing
```

Step 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.

Step 3. Add this VLAN to the desired interface. Run:

```
switch (config)# interface ethernet 1/1
switch (config ethernet 1/1)# switchport access vlan 20
```

Step 4. Create a VLAN interface. Run:

```
switch (config)# interface vlan 20
```

Step 5. Apply IP address to the VLAN interface.

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
```

Step 6. Enable the interface. Run:

```
switch (config interface vlan 20)# no shutdown
```

Configure VRRP:

This is the same configuration on both switches

Step 1. Enable VRRP protocol globally. Run:

```
switch (config)# protocol vrrp
```

Step 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
```

Step 3. Set the VIP address. Run:

```
switch (config interface vlan 20 vrrp 100)# address 20.20.20.40
```

Step 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
```

Step 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
```

Step 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
```

Step 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
```

Step 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.

6.6.3 Verifying VRRP

Step 1. Display VRRP brief status. Run:

```
switch(config)# show vrrp
Interface  VR  Pri  Time  Pre  State VR  IP addr
-----
Vlan20    1   200  2s    Y    Init      20.20.20.20
...
switch(config)#
```

Step 2. Display VRRP detailed status. Run:

```
switch (config)# show vrrp detail

VRRP Admin State : Enabled

Vlan20 - Group 1 (IPv4)

Instance Admin State : Enabled
State : Backup
Virtual IP Address : 20.20.20.40
Priority : 200
Advertisement interval (sec) : 2
Preemption : Enabled
Virtual MAC address : AA:BB:CC:DD:EE:FF
switch (config)#
```

Step 3. Display VRRP statistic counters. Run:

```
switch (config)# show vrrp statistics
Ethernet1/5 - Group 1 (IPv4)
Invalid packets:          0
Too short:                0
Transitions to Master     6
Total received:          155
Bad TTL:                  0
Failed authentication:    0
Unknown authentication:   0
Conflicting authentication: 0
Conflicting Advertise time: 0
Conflicting Addresses:    0
Received with zero priority: 3
Sent with zero priority:  3
switch (config)#
```

6.6.4 Commands

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 feature vrrp
Configuration Mode	Config
History	3.3.4500
Role	admin
Example	switch (config)# protocol vrrp
Related Commands	
Note	

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 is 1-255.
Default	N/A	
Configuration Mode	Config Interface VLAN	
History	3.3.4500	
Role	admin	
Example	switch (config interface vlan 10)# switch (config interface vlan 10 vrrp 10)#	
Related Commands		
Note	A maximum total of 255 VRRP instances are supported per switch system.	

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	
Role	admin	
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	
Related Commands		
Note	• 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 10 IP addresses associated with the VRRP instance. One primary and up to 10 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).	

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
Role	admin
Example	switch (config vrrp 100)# shutdown
Related Commands	
Note	

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 is 1-254.
Default	100	
Configuration Mode	Config VRRP Interface	
History	3.3.4500	
Role	admin	
Example	switch (config vrrp 100)# priority 200	
Related Commands		
Note	- The higher IP address will be selected as master, in case 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.	

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
Role	admin
Example	switch (config vrrp 100)# preempt
Related Commands	
Note	To set this router as backup for the current virtual router master, preempt must be enabled.

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. The password string must be up to 8 alphanumeric characters.
Default	Disabled	
Configuration Mode	Config VRRP Interface	
History	3.3.4500	
Role	admin	
Example	switch (config vrrp 100)# authentication text mypassword	
Related Commands		
Note		

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	
Role	admin	
Example	switch (config vrrp 100)# advertisement-interval 10	
Related Commands		
Note		

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	
Role	admin	
Example	<pre>switch(config)# show vrrp Interface VR Pri Time Pre State VR IP addr ----- Eth1/5 1 200 2s Y Init 192.0.1.10 ... switch(config)#</pre>	
Related Commands		
Note		

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	
Role	admin	
Example	<pre>switch (config)# show vrrp detail VRRP Admin State : Enabled Vlan20 - Group 1 (IPV4) Instance Admin State : Enabled State : Backup Virtual IP Address : 20.20.20.40 Priority : 200 Advertisement interval (sec) : 2 Preemption : Enabled Virtual MAC address : AA:BB:CC:DD:EE:FF switch (config)#</pre>	
Related Commands		
Note		

show vrrp statistics

show vrrp statistics [interface <type> <number>] [vr <id>]

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	
Role	admin	
Example	<pre>switch (config)# show vrrp statistics Ethernet1/5 - Group 1 (IPv4) Invalid packets: 0 Too short: 0 Transitions to Master 6 Total received: 155 Bad TTL: 0 Failed authentication: 0 Unknown authentication: 0 Conflicting authentication: 0 Conflicting Advertise time: 0 Conflicting Addresses: 0 Received with zero priority: 3 Sent with zero priority: 3 switch (config)#</pre>	
Related Commands		
Note		

6.7 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.



Designated traffic, such as ping to the MAGP interface is not supported. One of the two switches will be able to ping, so a ping from one switch can be done.

6.7.1 Configuring MAGP

Prerequisite steps:

Step 1. Enable IP routing functionality. Run:

```
switch (config)# ip routing
```

Step 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.

Step 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
```

Step 4. Create a VLAN interface. Run:

```
switch (config)# interface vlan 20
switch (config interface vlan 20)#
```

Step 5. Set an IP address to the VLAN interface. Run:

```
switch (config interface vlan 20)# ip address 11.11.11.11 /8
```

Step 6. Enable the interface. Run:

```
switch (config interface vlan 20)# no shutdown
```

➤ **To configure MAGP:**

Step 1. Enable MAGP protocol globally. Run:

```
switch (config)# protocol magp
```

Step 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.

Step 3. Set a virtual router primary IP address. Run:

```
switch (config interface vlan 20 magp 100)# ip virtual-router address 11.11.11.254
```



The IP address must be in the same subnet of the VLAN interface. This IP address is the default gateway for this MAGP instance. This should become the default gateway configured on the hosts connected to the relevant MLAG.

Step 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".

➤ **To verify the MAGP configuration, run:**

```
switch (config)# show magp 100  
MAGP 100  
  Interface vlan:20  
  MAGP state: Master  
  MAGP virtual IP: 11.11.11.254  
  MAGP virtual MAC: AA:BB:CC:DD:EE:FF  
switch (config)#
```



This output is to be expected in both MAGP switches.



For more advanced configuration options, please refer to the following Mellanox Community post: <https://community.mellanox.com/docs/DOC-1476>.

6.7.2 Commands

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
Role	admin
Example	switch (config)# protocol magp switch (config)#
Related Commands	
Note	IP routing must be enabled to enable MAGP.

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	
Role	admin	
Example	<pre>switch (config interface vlan 20)# magp 100 switch (config interface vlan 20 magp 100)#</pre>	
Related Commands		
Note	• Only one MAGP instance can be created on an interface • Different interfaces cannot share an MAGP instance • MAGP and VRRP are mutually exclusive	

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
Role	admin
Example	switch (config interface vlan 10 magp 1)# shutdown
Related Commands	
Note	

ip virtual-router address

ip virtual-router address <ip-address>
no ip virtual-router address

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.
Default	N/A	
Configuration Mode	Config Interface VLAN MAGP	
History	3.3.4500	
Role	admin	
Example	<pre>switch (config interface vlan 10 magp 1)# ip virtual-router address 10.10.10.10 switch (config interface vlan 10 magp 1)#</pre>	
Related Commands		
Note	The MAGP virtual IP address must be different from the interface IP address	

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	
Role	admin	
Example	switch (config interface vlan 10 magp 1)# ip virtual-router mac-address AA:BB:CC:DD:EE:FF switch (config interface vlan 10 magp 1)#	
Related Commands		
Note		

show magp

show magp [<instance> | 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	
Role	admin	
Example	<pre>switch (config)# show magp 3 Magp instance id: 3 Interface : vlan 10 Magp state: Active Magp virtual ip :192.168.1.1 Magp virtual MAC : 00:11:22:22:44:55 switch (config)#</pre>	
Related Commands		
Note		

6.8 DHCP Relay



DHCP Relay (DHCP-R) is not supported on SX10xx-xxxR and SX60xx-xxxR systems.

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 agents. These DHCP-R agents receive messages from DHCP clients and forward them to DHCP servers thus extending the reach of the DHCP beyond the local network.

6.8.1 DHCP-R 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 command [“vrf-auto-helper”](#) on page 1288 for more information.

6.8.2 Commands

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	
Role	admin	
Example	switch (config)# ip dhcp relay instance 1 switch (config ip dhcp relay instance 1)#	
Related Commands	N/A	
Note	If an instance is not specified then instance 1 is used (if nonexistent, then it is created).	

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
Role	admin	
Example	switch (config ip dhcp relay instance 1)# address 1.2.3.4	
Related Commands	ip dhcp relay	
Note	- 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).	

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
Role	admin	
Example	switch (config ip dhcp relay instance 1)# always-on	
Related Commands	ip dhcp relay	
Note	- 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).	

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
Role	admin
Example	switch (config ip dhcp relay instance 1)# information option
Related Commands	ip dhcp relay
Note	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).

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
Role	admin
Example	switch (config ip dhcp relay instance 1)# vrf 2
Related Commands	N/A
Note	• 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).

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	
Role	admin	
Example	switch (config ip dhcp relay instance 1)# port 65534	
Related Commands	N/A	
Note	- 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).	

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
Role	admin
Example	switch (config ip dhcp relay instance 1)# vrf-auto-helper
Related Commands	N/A
Note	• 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).

ip dhcp relay instance

ip dhcp relay [instance <instance-id>]
no ip dhcp relay [instance <instance-id>]

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
Default	N/A	
Configuration Mode	Config Interface VLAN Config Interface Port Channel Config Interface Ethernet set as router port interface	
History	3.6.3004	
Role	admin	
Example	switch (config interface vlan 10)# ip dhcp relay instance 7	
Related Commands	N/A	
Note	- In order to enable DHCP relay, at least one IP address should be configured, or always-on parameter should be turned on using the command “ip dhcp relay always-on” - When DHCP servers are configured, requests are forwarded only to configured servers - Only an existent DHCP-R may be specified	

clear ip dhcp relay counters

clear ip dhcp relay counters [vrf <vrf-name> | instance <instance-id>]

Clears all DHCP relay counters (all interfaces) in a given VRF or instance.

Syntax Description	vrf-name	VRF name
	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
Role	admin	
Example	switch (config)# clear ip dhcp relay counters	
Related Commands	N/A	
Note	• 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	

6.8.2.1 Interface

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
Role	admin	
Example	switch (config interface vlan 10)# ip dhcp relay information options circuit-id my-label	
Related Commands	N/A	
Note	The circuit ID sub-option is an IP interface attribute which is shared across all DHCP-R instances.	

6.8.2.2 Show

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 output and parameters
Role	admin	
Example	<pre>switch (config)# show ip dhcp relay instance 7 Instance ID: 7 VRF Name : default DHCP Servers: 1.1.2.1 DHCP relay agent options: always-on : Disabled Information Option : Disabled UDP port : 90 Auto-helper : Enabled Interface Label ----- vlan40 N/A vlan50 my-instance</pre>	
Related Commands	N/A	
Note	If no DHCP-R instance is given, then all DHCP-R instances are displayed	

show ip dhcp relay counters

show ip dhcp relay counters [vrf <vrf-name> | all]

Displays the DHCP relay counters in a given VRF.

Syntax Description	vrf	VRF name
	all	All VRF instances

Default	N/A				
Configuration Mode	Any Command Mode				
History	3.3.4150				
	3.6.1002	Added VRF and all parameters			
Role	admin				
Example	switch (config) # VRF Name: user				
	Interface	Received	Forwarded	Dropped	
	-----	-----	-----	-----	
	All Req	2	2	0	
	All Resp	2	2	0	
	Interface	Received	Forwarded	Dropped	Last Cleared
	-----	-----	-----	-----	-----
	Vlan10	2	2		0
	Vlan20	2	2		0
	VRF Name: default				
	Interface	Received	Forwarded	Dropped	
	-----	-----	-----	-----	
	All Req	3	2	1	
All Resp	3	3	0		
Interface	Received	Forwarded	Dropped	Last Cleared	
-----	-----	-----	-----	-----	
Vlan30	3	2		1	
Vlan40	3	3		0	
Related Commands	N/A				
Note					

7 InfiniBand Switching

7.1 Node Name

7.1.1 Commands

ib nodename

ib nodename <guid> name <name>
no ib nodename <guid>

Maps between GUID and node name.

Syntax Description	guid	The system GUID
	name	User defined string
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib nodename 00:00:00:00:60:04:03:30 name my-name switch (config) # show ib nodename GUID='00:00:00:00:60:04:03:30', name='my-name', discovered='no' switch (config) #</pre>	
Related Commands		
Notes	If an entry with GUID exists, the existing name will be replaced with a new name.	

show ib nodename

show ib nodename

Maps between GUID and node name.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.1.0000
Role	admin
Example	<pre>switch (config) # show ib nodename GUID='00:00:00:00:60:04:03:30', name='my-name', discovered='no' switch (config) #</pre>
Related Commands	ib nodename
Notes	

7.2 Fabric

7.2.1 Commands

fabric zero-counters

fabric zero-counters

Clears the performance counters of the node.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	monitor/admin
Example	switch (config) # fabric zero-counters Counters zeroed successfully switch (config) #
Related Commands	
Notes	

show fabric

show fabric {pm | sm}

Displays InfiniBand fabric details.

Syntax Description	pm	Displays InfiniBand fabric performance measurements.
	sm	Displays InfiniBand fabric SMs.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	First version
	3.4.0000	Added note
Role	admin	
Example	<pre>switch (config) # show fabric sm % # This database file was automatically generated by IBDIAG ibdiagnet fabric SM report SM - master Port=0 lid=0x0005 guid=0x0002c903004a2980 dev=51000 priority:15 SM - standby Port=0 lid=0x0001 guid=0x0000000000000111 dev=51000 priority:0 switch (config) #</pre>	
Related Commands		
Notes	This command requires a fabric inspector license (LIC-fabric-inspector).	

show guides

show guides

Displays GUIDs per ASIC in the chassis.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000 3.4.2008 Updated Example 3.6.1002 Updated Example
Role	admin
Example	<pre>switch (config) # show guides ===== Module Device IB Subnet GUID ===== SYSTEM - - E4:1D:2D:03:00:2E:49:40 MGMT SIB infiniband-default E4:1D:2D:03:00:2E:49:40 MGMT SIB infiniband-1 E4:1D:2D:03:00:2E:49:41 MGMT SIB infiniband-2 E4:1D:2D:03:00:2E:49:42 switch (config) #</pre>
Related Commands	
Notes	

show system guid

show {guids | system guid}

Displays the system GUID.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Role	admin
Example	switch (config) # show system guid 00:02:C9:03:00:43:D9:00 switch (config) #
Related Commands	
Notes	

show lids

show lids

Displays the LIDs of each module in the switch system

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000 3.4.2008 Updated Example 3.6.1002 Updated Example
Role	admin/monitor
Example	<pre>switch (config) # show lids ===== Module Device IB Subnet LID ===== MGMT SIB infiniband-default 1 MGMT SIB infiniband-1 8 MGMT SIB infiniband-2 3 switch (config) #</pre>
Related Commands	
Notes	

7.3 IB Router

IB router provides the ability to send traffic between two or more IB subnets thereby potentially expanding the size of the network to over 40k end-ports, enabling separation and fault resilience between islands and IB subnets, and enabling connection to different topologies used by different subnets.

The forwarding between the InfiniBand subnets is performed using GRH (global route header) lookup.

IB router capabilities are supported only on SB7780 switch system which comes with the following default configuration:


- L3 capabilities enabled
- 2 SWIDs, with interface 1/1 mapped to infiniband-default and interface 1/2 mapped to infiniband-1

The IB router's basic functionality includes:

- Removal of current L2 LRH (local routing header)
- Routing table lookup – using GID from GRH
- Building new LRH according to the destination and the routing table


The DLID in the new LRH is built using simplified GID-to-LID mapping (where LID = 16 LSB bits of GID) thereby not requiring to send for ARP query/lookup.

Figure 42: Site-Local Unicast GID Format



For this to work, the SM allocates an alias GID for each host in the fabric where the alias GID = {subnet prefix[127:64], reserved[63:16], LID[15:0]}. Hosts should use alias GIDs in order to transmit traffic to peers on remote subnets.

Figure 43: Host-to-Host IB Router Unicast Flow



For more information on IB router architecture and functionality, please refer to the following Mellanox Community page: <https://community.mellanox.com/docs/DOC-2384>.



IB router requires HCA configuration such as SM, partition key, MPI, GID translation, and more. To learn more about these configurations, please refer to the following Mellanox Community page: <https://community.mellanox.com/docs/DOC-2466>.

7.3.1 Configuring IB Router

Prerequisites:

Step 1. Check system capabilities to make sure IB L3 is supported. Run:

```
switch (config) # show system capabilities
IB: Supported, L2, L3, Adaptive Routing
Max SM nodes: 2048
IB Max licensed speed: EDR
```

Step 2. Configure system profile to multi-switch with 2 SWIDs. Run:

```
switch (config) # system profile ib num-of-swids 2 ib-router
```



Note that some of the interfaces may not be mapped to a SWID.

Step 3. Verify system profile configuration. Run:

```
switch (config) # show system profile
Profile:          ib
Number of SWIDs:  2
Adaptive Routing: no
IB Routing:       yes
```

➤ **To configure IB router:**

Step 1. Map an interface to a SWID. Run:

```
switch (config) # interface ib 1/1 switchport access subnet infiniband-default force
switch (config) # interface ib 1/2 switchport access subnet infiniband-1 force
```

Step 2. Verify SWID configuration. Run:

```
switch (config) # show interfaces ib status
```

Interface	Description	IB Subnet	Speed	Current line rate	Logical port state	Physical port state
IB1/1		infiniband-default	-	-	Down	Polling
IB1/2		infiniband-1	edr	100.0 Gbps	Initialize	LinkUp
IB1/3		-	-	-	-	-
...						

Step 3. Configure and enable IB router. Run:

```
switch (config) # ib router
switch (config) # no ib router shutdown
```

Step 4. Enable IB subnet interface. Run:

```
switch (config) # no interface ib-subnet infiniband-default shutdown
switch (config) # no interface ib-subnet infiniband-1 shutdown
```

Step 5. Verify configuration. Run:

```
switch (config) # show ib router
Routing state: enabled

IB subnet          Routing enabled
infiniband-default enabled
infiniband-1       enabled
```

```
switch (config) # show interfaces ib-subnet infiniband-default
infiniband-default state:
  GUID           : F4:52:14:03:00:6E:F2:8B
  Alias GUID     : N/A
  LID            : 10
  Subnet prefix  : FE:C0:00:00:00:00:00:08
  Physical state : LinkUp
  Logical state  : Active
  L3 interface state : Up
switch (config) #
```



For more advanced information on IB router configuration, please refer to the following Mellanox Community page: <https://community.mellanox.com/docs/DOC-2466>.

7.3.2 Subnet Prefix Checking

The SB7780 IB router expects the subnet prefix to be constructed according to some very specific rules. By default, the command which enables IB routers validates the subnet prefix prior to allowing the change.

The commands which affect subnet prefix checking are as follows:

- `ib sm <name> enable` – starts SM on this node or any node in cluster
- `ib sm subnet-prefix <subnet-prefix>` – configures the subnet prefix
- `ib sm rtr-aguid-enable <1 | 2>` – enables support for alias GIDs as needed by IB routers

When any of these commands is run, while the other two have already been issued, the value of the subnet prefix is checked. If it is not valid, the current commit is rejected and the OpenSM state does not change.

➤ *To disable subnet prefix checking*

Step 1. Verify the status of subnet prefix override. Run:

```
switch (config) # show ib sm subnet-prefix-override
enable
```

Step 2. If enabled, disable subnet-prefix-override. Run:

```
switch (config) # ib sm subnet-prefix-override
```

Step 3. Verify configuration. Run:

```
switch (config) # show ib sm subnet-prefix-override
disable
```

7.3.3 Commands

ib router

ib router
no ib router

Enables the set of commands that allow control of IB router functionality.
The no form of the command disables IB router commands and removes all related configurations.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.6.0500
Role	admin
Example	switch (config) # ib router switch (config) #
Related Commands	“system profile” on page 240
Notes	

ib router shutdown

ib router shutdown
no ib router shutdown

Disables IB router.
The no form of the command enables IB router.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config
History	3.6.0500
Role	admin
Example	switch (config) # no ib router shutdown switch (config) #
Related Commands	
Notes	This command does not clear IB router configuration

interface ib-subnet

interface ib-subnet <swid-name>
no interface ib-subnet <swid-name>

Creates routing on IB router subnet.
The no form of the command removes routing on router interface.

Syntax Description	swid-name	Name of the SWID: infiniband-default, infiniband-1...infiniband-5
Default	N/A	
Configuration Mode	Config	
History	3.6.0500	
Role	admin	
Example	switch (config) # interface ib-subnet infiniband-3 switch (config) #	
Related Commands	system profile	
Notes	The maximum number of SWIDs depends on the number of SWIDs defined in the profile	

interface ib-subnet shutdown

interface ib-subnet <swid-name> shutdown
no interface ib-subnet <swid-name> shutdown

Disables routing on IB router subnet.
 The no form of the command enables routing on router interface.

Syntax Description	swid-name	Name of the SWID: infiniband-default, infiniband-1...infiniband-5
	shutdown	Admin down on router interface Admin up on router interface with no form of command
Default	Disabled	
Configuration Mode	Config	
History	3.6.0500	
Role	admin	
Example	switch (config) # no interface ib-subnet infiniband-3 shutdown switch (config) #	
Related Commands		
Notes		

show ib router

show ib router

Displays current IB router functionality.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.6.0500
Role	admin
Example	<pre>switch (config) # show ib router Routing state: enabled IB Subnet Routing enabled infiniband-default enabled infiniband-1 disabled infiniband-2 enabled infiniband-3 enabled switch (config) #</pre>
Related Commands	
Notes	

show interfaces ib-subnet

show interfaces ib-subnet [<swid-name>] [brief]

Displays statistics of one or all IB subnets with enabled IB routing.

Syntax Description	swid-name	Name of the SWID: infiniband-default, infiniband-1...infiniband-5
	brief	Displays output in a table format
Default	Disabled	
Configuration Mode	Config	
History	3.6.0500	
Role	admin	
Example	<pre>switch (config) # show interfaces ib-subnet infiniband-3 infiniband-3 state: GUID : F4:52:14:03:00:6E:F2:8B Alias GID : N/A LID : 10 Subnet prefix : FE:C0:00:00:00:00:00:08 Physical state : LinkUp Logical state : Active L3 interface state : Up</pre>	
Related Commands		
Notes		

7.4 Interface

7.4.1 Transceiver Information

MLNX-OS® 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 ib 1/36 transceiver
Slot 1 port 36 state
  identifier           : QSFP+
  cable/module type    : Passive copper, unequalized
  infiniband speeds    : SDR , DDR , QDR , FDR
  vendor              : Mellanox
  cable length         : 2m
  part number          : MC2207130-0A1
  revision            : A3
  serial number        : MT1324VS02215

switch (config) #
```



The indicated cable length is rounded up to the nearest natural number.

7.4.2 High Power Transceivers

Mellanox switch systems offer high power transceiver (e.g. LR4) support in the following ports:

- SX6036 – ports 1, 3, 33, 35
- SX6012/SX6710/SX6720 – all ports
- SB7700 – all ports
- TX6000/TX6100 – all uplink ports

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 ib” is run.

7.4.3 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 of errors that can be corrected depends on the FEC algorithm.



Switch-IB™ EDR (100Gb/s) Mellanox-to-Mellanox InfiniBand connections enable standard low-latency Reed Solomon (LL RS) FEC on active optical cables longer than 30 meters and passive copper cables longer than 2m.

7.4.4 Commands

interface ib

interface ib [internal] {<inf> <inf-range>}		
Enters the InfiniBand interface configuration mode.		
Syntax Description	[internal] <inf>	For 1U switches: interface 1/<interface> For director switches: interface ib <interface> interface ib internal leaf <interface> interface ib internal spine <interface>
	inf-range	Enters the configuration mode of a range of interfaces. Format: <slot>/<port>-<slot>/<port>
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
	3.4.2008	Added internal leaf and spine options
Role	admin	
Example	switch (config) # interface ib 1/1 switch (config interface ib 1/1) #	
Related Commands	show interfaces ib	
Notes	Interface range (inf-range) option is not valid on director switch systems.	

mtu

mtu <frame-size>

Configures the Maximum Transmission Unit (MTU) frame size for the interface.

Syntax Description	frame-size	Possible Value for MTU	
		• 256	256 bytes
		• 512	512 bytes
		• 1K	1K bytes
		• 2K	2K bytes
		• 4K	4K bytes
Default	4096 bytes		
Configuration Mode	Config Interface IB		
History	3.1.0000		
Role	admin		
Example	switch (config interface ib 1/1) # mtu 4K switch (config interface ib 1/1) #		
Related Commands	show interfaces ib		
Notes			

shutdown

shutdown
no shutdown

Disables the interface.
The no form of the command enables the interface.

Syntax Description	N/A
Default	The interface is enabled.
Configuration Mode	Config Interface IB
History	3.1.0000
Role	admin
Example	switch (config interface ib 1/1) # shutdown switch (config interface ib 1/1) #
Related Commands	show interfaces ib
Notes	N/A

description

description <string>

Sets an interface description.

Syntax Description	string	40 bytes
Default	""	
Configuration Mode	Config Interface IB	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config interface ib 1/1) # description my-interface switch (config interface ib 1/1) #</pre>	
Related Commands	show interfaces ib	
Notes		

speed

speed <port speed> [force]

Sets the speed negotiation of the interface.

Syntax Description	port speed	The following options are available: • sdr – 10.0Gb/s rate on 4 lane width • ddr – 20.0Gb/s rate on 4 lane width • qdr – 40.0Gb/s rate on 4 lane width • fdr10 – 40.0Gb/s rate on 4 lane width • fdr – 56.0Gb/s rate on 4 lane width • edr – 100.0Gb/s rate on 4 lane width
	force	Forces configuration of speed-list not containing SDR bit
Default	Depends on the port module type, not all interfaces support all speed options	
Configuration Mode	Config Interface IB	
History	3.1.0000	
	3.4.1604	Updated Syntax Description and Example
Role	admin	
Example	<pre>switch (config interface ib 1/1) # speed fdr10 fdr edr switch (config interface ib 1/1) #</pre>	
Related Commands	show interfaces ib	
Notes	• This command is backwards compatible so old configuration file containing this command with the old form (with legal bit mask) are still supported • Configuring more than one speed is possible by typing in consecutive speed names separated by spaces • If the speed-options list does not include SDR speed, it is configured automatically. However, if the force option is used (supported on FDR10 only), SDR is not configured. • If the other side of the link is a SwitchX® or ConnectX®-3 device, to allow the link to raise in FDR speed, QDR speed must also be allowed	

op-vls

op-vls <value>

Sets the operational VLs of the interface.

The no form of the command sets the operational VLs to its default value.

Syntax Description	value	Possible value for operational VLs
		• 1 VL0 • 2 VL0, VL1 • 4 VL0 - VL3 • 8 VL0 - VL7
Default	8 (VL0 - VL7)	
Configuration Mode	Config Interface IB	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config interface ib 1/1) # op-vls 1 switch (config interface ib 1/1) #</pre>	
Related Commands	show interfaces ib	
Notes		

width

width <value>

Sets the width of the interface.

The no form of the command sets the speed of the interface to its default value.

Syntax Description	value	Possible value for width: • 1 – 1X • 5 – 1X, 4X
Default	5 (1X, 4X)	
Configuration Mode	Config Interface IB	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config interface ib 1/1) # width 1 switch (config interface ib 1/1) #</pre>	
Related Commands	show interfaces ib	
Notes		

clear counters

clear counters

Clears the interface counters.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config Interface IB
History	3.1.0000
Role	admin
Example	switch (config interface ib 1/1) # clear counters switch (config interface ib 1/1) #
Related Commands	show interfaces ib
Notes	

interface ib internal notification link-speed-mismatch

interface ib internal notification link-speed-mismatch [<time>]
no interface ib internal notification link-speed-mismatch

Enables notifications on internal link speed mismatch in SNMP.
The no form of the command disables notifications on internal links speed mismatch in SNMP.

Syntax Description	time	Enables periodic notifications (traps and log) on internal link speed mismatch status. The time is in hours. "0" disables the feature
Default	Disabled	
Configuration Mode	Config	
History	3.4.3000	
Role	admin	
Example	<pre>switch (config) # interface ib internal link-speed-mismatch 6 switch (config) # show interfaces ib internal notification ===== Internal links information ===== State change enabled : no Speed mismatch enabled : yes Periodic notifications : 6 (hours) switch (config) #</pre>	
Related Commands	show interfaces ib internal notification	
Notes	Link-speed-mismatch shows internal link entries in the ifVPITable	

interfaces ib internal notification link-state-change

interfaces ib internal notification link-state-change
no interfaces ib internal notification link-state-change

Enables notifications on internal links state change in SNMP.
 The no form of the command disables notifications on internal links state change in SNMP.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config
History	3.3.4318 3.3.4550 Added note
Role	admin
Example	switch [master] (config) # interfaces ib internal notification switch [master] (config) #
Related Commands	show interfaces ib internal notification
Notes	Link-state-change shows internal link entries in the ifTable and the ifXTable

switchport access subnet

switchport access subnet <swid-name> [force]
no switchport access subnet <swid-name> [force]

Maps interface to SWID.

The no form of the command unmaps an interface from a SWID.

Syntax Description	swid-name	Name of the SWID: infinibad-default, infiniband-1...infinibad-5
	force	Forces configuration (no need to shutdown interface before running command)
Default	Unmapped	
Configuration Mode	Config Interface IB	
History	3.6.0500	
Role	admin	
Example	switch (config interface ib1/36) # switchport access subnet infiniband-1	
Related Commands		
Notes	• Mapping an interface automatically enables it • Remapping an interface resets all its configuration except for interface description • Unmapping an interface resets all its configuration except for interface description • An interface needs to be disabled before remapping/unmapping unless the “force” parameter is used	

show interfaces ib

show interfaces ib <inf>

Displays the configuration and status for the interface.

Syntax Description	internal	Internal interfaces.
	inf	- Slot/Port (i.e. 1/1) - LXX/SXX (i.1 L01 or S01)
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.1.0000	
	3.4.1604	Updated Example
	3.6.1002	Updated Example
Role	admin	
Example	<pre>switch (config) # show interfaces ib 1/1 IB1/1 state: Logical port state : Down Physical port state : Polling Current line rate : - Supported speeds : sdr, ddr, qdr, fdr, edr Speed : - Supported widths : 1X, 4X Width : 4X Max supported MTUs : 4096 MTU : 256 VL capabilities : VL0 - VL7 Operational VLS : VL0 - VL7 Description : IB Subnet : infiniband-default Phy-profile : high-speed-ber Width reduction mode : RX bytes : 0 RX packets : 0 RX errors : 0 Symbol errors : 0 VL15 dropped packets : 0 TX bytes : 0 TX packets : 0 TX wait : 0 TX discarded packets : 0 switch (config) #</pre>	

Related Commands

Notes

If a high power transceiver (e.g. LR4) is inserted to a port that does not support it, the link will not go up, and the following warning message is displayed: “Warning: High power transceiver is not supported” when the command “show interfaces ib” is run. For more information, please refer to [Section 7.4.2, “High Power Transceivers,” on page 1311](#).

show interfaces ib status

show interfaces ib [<inf>] status

Displays the status, speed and negotiation mode of the specified interface.

Syntax Description	internal	Internal interfaces					
	leaf-ports	filter to leaf-ports only					
	inf	Interface number: <slot>/<port>					
Default	N/A						
Configuration Mode	Any Command Mode						
History	3.2.0500						
	3.4.1604	Updated Example					
	3.6.1002	Updated Example					
Role	admin						
Example	switch (config) # show interfaces ib status						
	Interface	Description	IB Subnet	Speed	Current line rate	Logical port state	Physical port state
	-----	-----	-----	-----	-----	-----	-----
	IB1/1		infiniband-1	fdr	56.0 Gbps	Active	LinkUp
	IB1/2		infiniband-2	fdr	56.0 Gbps	Active	LinkUp
	IB1/3		infiniband-default	-	-	Down	Polling
	IB1/4		infiniband-default	-	-	Down	Polling
	IB1/5		infiniband-default	-	-	Down	Polling
	IB1/6		infiniband-default	-	-	Down	Polling
	IB1/7		infiniband-default	-	-	Down	Polling
	IB1/8		infiniband-default	-	-	Down	Polling
	IB1/9		infiniband-default	-	-	Down	Polling
	IB1/10		infiniband-default	-	-	Down	Polling
	IB1/11		infiniband-default	-	-	Down	Polling
							
	switch (config) #						
Related Commands							
Notes							

show interfaces ib internal

show interfaces ib internal [leaf | spine] [<slot/module/port>]

Displays running state for the internal ports of leafs or spines.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.2.0500
Role	admin
Example	<pre>switch (config) # show interfaces ib internal spine 1/1/4 IB1/1/4 state: Connected to slot/chip : 4/1 Connected to port : 19 Connected device active: 1 Error state : 0 Logical port state : Active Physical port state : LinkUp Current line rate : 56.0 Gbps Supported speeds : sdr, ddr, qdr, fdr10, fdr Speed : fdr Supported widths : 1X, 4X Width : 4X Max supported MTUs : 4096 MTU : 4096 VL capabilities : VL0 - VL7 Operational VLS : VL0 - VL7 Description : Phy-profile : high-speed-ber Width reduction mode : disabled</pre> switch (config) #

Related Commands

Notes

show interfaces ib internal capabilities

show interfaces ib internal [leaf | spine] [<slot/module/port>] capabilities

Displays capabilities of internal leaf or spine interfaces.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.2.0500
Role	admin
Example	<pre>switch (config) # show interfaces ib internal leaf 1/1/26 capabilities IB1/1/26 LLR: FDR10, FDR, switch (config) #</pre>
Related Commands	
Notes	

show interfaces ib internal llr

show interfaces ib internal [leaf | spine] [<slot/module/port>] llr

Displays LLR state of internal leaf or spine interfaces.

Syntax Description	N/A		
Default	N/A		
Configuration Mode	Any Command Mode		
History	3.2.0500		
Role	admin		
Example	switch (config) # show interfaces ib internal leaf 1/1/26 llr		
	Interface	phy-profile	LLR status
	IB1/1/26	high-speed-ber	Active
	switch (config) #		
Related Commands			
Notes			

show interfaces ib internal status

show interfaces ib internal [leaf | spine] [<slot/module/port>] status

Displays detailed running state of internal leaf or spine interfaces.

Syntax Description	N/A																							
Default	N/A																							
Configuration Mode	Any Command Mode																							
History	3.2.0500																							
Role	admin																							
Example	<pre>switch (config) # show interfaces ib internal leaf 1/1/26 status</pre> <table><thead><tr><th>Interface state</th><th>Description</th><th>Speed</th><th>Current line rate</th><th>Logical port state</th><th>Physical port</th></tr><tr><th>-----</th><th>-----</th><th>-----</th><th>-----</th><th>-----</th><th>-----</th></tr></thead><tbody><tr><td>IB1/1/26</td><td></td><td>fdr</td><td>56.0 Gbps</td><td>Active</td><td>LinkUp</td></tr></tbody></table> <pre>switch (config) #</pre>						Interface state	Description	Speed	Current line rate	Logical port state	Physical port	-----	-----	-----	-----	-----	-----	IB1/1/26		fdr	56.0 Gbps	Active	LinkUp
Interface state	Description	Speed	Current line rate	Logical port state	Physical port																			
-----	-----	-----	-----	-----	-----																			
IB1/1/26		fdr	56.0 Gbps	Active	LinkUp																			
Related Commands																								
Notes																								

show interfaces ib transceiver

show interfaces ib [<inf>] transceiver

Displays the transceiver info.

Syntax Description	inf interface number: <slot>/<port>
Default	N/A
Configuration Mode	Any Command Mode
History	3.1.0000
Role	admin
Example	<pre>switch (config) # show interfaces ib 1/1 transceiver Slot L01 port 13 state identifier : QSFP+ cable/module type : Passive copper, unequalized infiniband speeds : SDR , DDR , QDR vendor : Mellanox cable length : 2 m part number : MC2207130-002 revision : B0 serial number : AA051150077 switch (config) #</pre>
Related Commands	
Notes	- For a full list of the supported cables and transceivers, please refer to the LinkX™ Cables and Transceivers webpage in Mellanox.com: http://www.mellanox.com/page/cables?mtag=cable_overview. - If a high power transceiver (e.g. LR4) is used, it will be indicated in the field “cable/module type”

show interfaces ib transceiver diagnostics

show interfaces ib [<inf>] transceiver diagnostics

Displays cable channel monitoring and diagnostics info for this interface.

Syntax Description	inf	Interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.6.2002	
Role	admin	

Example

```
switch (config) # show interfaces ib 1/1 transceiver diagnostics
Port 1/1 transceiver diagnostic data:
```

```

Temperature (-127C to +127C)
  Temperature           : 33 C
  Hi Temp Alarm Thresh  : 17 C
  Low Temp Alarm Thresh : 2 C
  Temperature Alarm     : None

Voltage ( 0 to 6.5535 V)
  Voltage               : 3.29450 V
  Hi Volt Alarm Thresh  : 3.70000 V
  Low Volt Alarm Thresh : 2.90000 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 to 6.5535 mW)
  Ch1 Tx Power          : 1.03080 mW
  Ch2 Tx Power          : 1.05070 mW
  Ch3 Tx Power          : 1.07150 mW
  Ch4 Tx Power          : 1.10180 mW
  Hi Tx Power Alarm Thresh : 3.46730 mW
  Low Tx Power Alarm Thresh : 0.07240 mW
  Ch1 Tx Power Alarm     : None
  Ch2 Tx Power Alarm     : None
  Ch3 Tx Power Alarm     : None
  Ch4 Tx Power Alarm     : None

Rx Power ( 0 to 6.5535 mW)
  Ch1 Rx Power          : 1.13980 mW
  Ch2 Rx Power          : 1.11720 mW
  Ch3 Rx Power          : 1.08800 mW
  Ch4 Rx Power          : 1.16450 mW
  Hi Rx Power Alarm Thresh : 0.33000 mW
  Low Rx Power Alarm Thresh : 1.01830 mW
  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) : 12-05-2016
```

Related Commands

Note

This example is for a QSFP transceiver

show interfaces ib transceiver raw

show interfaces ib [<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	
Role	admin	
Example	<pre>switch (config) # show interfaces ib 1/7 transceiver raw IB1/7 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 0010 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 0030 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 0050 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 0070 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 00 05 8d 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 00f0 00 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 0090 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 00b0 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 00d0 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 00f0 00 00 00 00 00 00 00 00 00 00 00 00 00 00</pre>	
Related Commands		
Notes		

show interfaces ib capabilities

show interfaces ib <inf> capabilities

Shows interface capabilities.

Syntax Description	inf	Slot/port (i.e. 1/1).
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.2.0500	
Role	admin	
Example	switch (config) # show interfaces ib 1/1 capabilities Ib 1/1 LLR: FDR10, FDR, switch (config)	
Related Commands		
Notes		

7.5 Subnet Manager (SM)

The InfiniBand Subnet Manager (SM) is a centralized entity running in the switch. The SM discovers and configures all the InfiniBand fabric devices to enable traffic flow between those devices.

The SM applies network traffic related configurations such as Quality of Service (QoS), routing, and partitioning of the fabric devices. You can view and configure the Subnet Parameters (SM) via the CLI/WebUI menu. The embedded SM on the MLNX-OS can be used to manage fabrics up to 648 nodes.

7.5.1 Enabling Subnet Manager

The SM is used to discover and configure all the InfiniBand fabric devices to enable traffic flow between those devices.

➤ **To enable Subnet Manager:**

Step 1. Enable Subnet Manager (disabled by default). Run:

```
switch (config) # ib smnode my-sm enable
```

Step 2. (Optional) Set the priority for the Subnet Manager. Run:

```
switch (config) # ib smnode my-sm sm-priority <priority>
```

7.5.2 Partitions

Partitioning enforces isolation among systems sharing an InfiniBand fabric. Partitioning is not related to boundaries established by subnets, switches, or routers. Rather, a partition describes a set of end nodes within the fabric that may communicate. Each port of an end node is a member of at least one partition and may be a member of multiple partitions. A partition manager (part of the SM) assigns partition keys (PKEYs) to each channel adapter port. Each PKEY represents a partition. Reception of an invalid PKEY causes the packet to be discarded. Switches and routers may optionally be used to enforce partitioning. In this case the partition manager programs the switch or router with PKEY information and when the switch or router detects a packet with an invalid PKEY, it discards the packet.

Fabric administration can assign certain Service Levels (SLs) for particular partitions. This allows the SM to isolate traffic flows between those partitions, and even if both partitions operate at the same QoS level, each partition can be guaranteed its fair share of bandwidth regardless of whether nodes in other partitions misbehave or are over subscribed.

The switch enables the configuration of partitions in an InfiniBand fabric.

The default partition is created by the SM unconditionally (whether it was defined or not).

7.5.2.1 Relationship with ib0 Interface

IP interface “ib0” is running under the default PKEY (0x7fff) and can be used for in-band management connectivity to the system.

7.5.2.2 Configuring Partition



The partitions configuration is applicable and to be used only when the SM is enabled and running on the system.

➤ **To configure a partition:**

Step 1. Create a partition. Run:

```
switch (config) # ib partition my-partition pkey 0x7ff2
```

Step 2. Enter partition configuration mode. Run:

```
switch (config) # partition my-partition
switch (config partition name my-partition) #
```

Step 3. Add partition members. Run:

```
switch (config partition my-partition) # member all
```

Step 4. Verify the partition configuration. Run:

```
switch (config partition my-partition) # show ib partition
Default
  PKey      = 0x7FFF
  defmember = full
  ipoib     = yes
members
  GUID='ALL' member='full'
my-partition
  PKey      = 0x7ff2
members
  GUID='ALL' member='default'
switch (config partition name my-partition) #
```

7.5.3 Adaptive Routing

Adaptive routing (AR) allows optimizing data traffic flow. The InfiniBand protocol uses multiple paths between any two points. Thus, when unexpected traffic patterns cause some paths to be overloaded, AR can automatically move traffic to less congested paths according to the current temporal state of the network.



The embedded SM over the switch does not support configuring adaptive routing. To use this option in the fabric please use an external SM.

AR support is enabled by default on system profile “ib-single-switch”. To disable AR run the command system profile ib-no-adaptive-routing-single-switch.



The AR option needs to be enabled in the SM for it to take affect.

7.5.4 Commands

7.5.4.1 Subnet Manager (SM)

ib sm

ib sm

no ib sm

Enables the SM on this node.

The no form of the command disables the SM on this node.

Syntax Description	N/A
Default	disable
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # ib sm switch (config) #
Related Commands	show ib sm
Notes	

ib sm accum-log-file

ib sm accum-log-file
no ib sm accum-log-file

Adds SM log entries at the end of the current log.
The no form of the command overwrites SM log file on every restart.

Syntax Description	N/A
Default	Enabled
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # ib sm accum-log-file switch (config) # show ib sm accum-log-file enable switch (config) #</pre>
Related Commands	show ib sm accum-log-file
Notes	

ib sm allow-both-pkeys

ib sm allow-both-pkeys
no ib sm allow-both-pkeys

Enables having both full and limited membership on the same partition.
The no form of the command disables having both full and limited membership on the same partition.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config
History	3.4.1100
Role	admin
Example	switch (config) # ib sm allow-both-pkeys switch (config) #
Related Commands	defmember member
Notes	

ib sm babbling-policy

ib sm babbling-policy
no ib sm babbling-policy

Enables the SM to disable babbling ports (i.e., generating frequent traps).
The no form of the command disables the SM babbling policy.

Syntax Description	N/A
Default	disable
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # no ib sm babbling-policy switch (config) # show ib sm babbling-policy disable switch (config) #</pre>
Related Commands	show ib sm babbling-policy
Notes	In case the babbling policy is enabled, and decides to close a babbling interface (one which sends 129,130,131 traps, for example), the SM disables the port.

ib sm connect-roots

ib sm connect-roots
no ib sm connect-roots

Forces the routing engine to make connectivity between root switches.
 The no form of the command disables logical LID path between root switches.

Syntax Description	N/A
Default	true
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # ib sm connect-roots switch (config) # show ib sm connect-roots true switch (config) #</pre>
Related Commands	show ib sm connect-roots
Notes	• This command is relevant only for 'updn' and 'free' algorithm (refer to 'ib sm routing-engines' command) • This option enforces routing engines (up/down and fat-tree) to make connectivity between root switches and in this way to be fully IBA complaint. This may violate the "deadlock-free" status of the algorithm. Hence, it is recommended to use the command carefully.

ib sm drop-event-subscription

ib sm drop-event-subscription
no ib sm drop-event-subscription

Configures IB SM to drop interface subscribe or unsubscribe events.
The no form of the command resets this parameter to its default value.

Syntax Description	N/A
Default	IB SM does not drop interface subscribe or unsubscribe events
Configuration Mode	Config
History	3.4.2008
Role	admin
Example	switch (config) # ib sm drop-event-subscription switch (config) #
Related Commands	
Notes	

ib sm enable-quirks

ib sm enable-quirks
no ib sm enable-quirks

Enables the SM to use high risk features and handle hardware workarounds. The no form of the command disables the SM from using high risk features and hardware workarounds.

Syntax Description	N/A
Default	disable
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # ib sm enable-quirks switch (config) # show ib sm enable-quirks enable switch (config) #</pre>
Related Commands	show ib sm enable-quirks
Notes	

ib sm exit-on-fatal

ib sm exit-on-fatal
no ib sm exit-on-fatal

Enables the SM to exit upon fatal initialization errors.
The no form of the command disables the SM from exiting upon fatal initialization errors.

Syntax Description	N/A
Default	enable
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # ib sm exit-on-fatal switch (config) # show ib sm exit-on-fatal enable switch (config) #</pre>
Related Commands	show ib sm exit-on-fatal
Notes	

ib sm force-link-speed

ib sm force-link-speed <speed-options>
no ib sm force-link-speed

Defines the SM behavior for PortInfo:LinkSpeedEnabled, PortInfo:LinkSpeedExtEnabled and MLNX ExtendedPortInfo on the switch ports.

Syntax Description	speed-options	The following options are available: • sdr – 10.0 Gb/s rate on 4 lane width • ddr – 20.0 Gb/s rate on 4 lane width • qdr – 40.0 Gb/s rate on 4 lane width • fdr10 – 40.0 Gb/s rate on 4 lane width • fdr – 56.0 Gb/s rate on 4 lane width • edr – 100.0 Gb/s rate on 4 lane width
Default	Set to PortInfo:LinkSpeedExtSupported	
Configuration Mode	Config	
History	3.1.0000 3.4.1604 Updated Syntax Description, Example and Notes	
Role	admin	
Example	<pre>switch (config) # ib sm force-link-speed sdr ddr qdr fdr10 switch (config) #</pre>	
Related Commands	<pre>show ib sm force-link-speed show ib sm force-link-speed-ext show ib sm fdr10</pre>	
Notes	• The following options, as defined in <i>InfiniBand Specification 1.2.1</i> section 14.2.5.6, table 145 “PortInfo” • This command updates force-link-speed, force-link-speed ext and fdr10 which are open sm parameters • This command is backwards compatible so old configuration file containing this command with the old form (with legal bit mask) are still supported • If the speed-options list does not include SDR speed, it is configured automatically • Configuring more than one speed is possible by typing in consecutive speed names separated by spaces	

ib sm force-log-flush

ib sm force-log-flush
no ib sm force-log-flush

Forces every log message generated to be flushed.
The no form of the command does not force a flush after every log write.

Syntax Description	N/A
Default	disable
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # ib sm force-log-flush switch (config) # show ib sm force-log-flush enable switch (config) #</pre>
Related Commands	show ib sm force-log-flush
Notes	

ib sm guid2lid-cache

ib sm guid2lid-cache
no ib sm guid2lid-cache

Allows SM to use cached GUID-to-lid mapping data. When enabled, the SM honors the cached GUID-to-lid mapping information if:

- It exists
- It is valid
- sm_reassign_lids is disabled

The no form of the command disallows use of cached GUID-to-lid mapping data.

Syntax Description	N/A
Default	disable
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # ib sm guid2lid-cache switch (config) # show ib sm guid2lid-cache enable switch (config) #</pre>
Related Commands	show ib sm guid2lid-cache
Notes	

ib sm honor-partitions

ib sm honor-partitions
no ib sm honor-partitions

Sets the no_partition_enforcement flag to 0. This setting controls global support for partitioning in the subnet.

The no form of the command disables subnet partition support.

Syntax Description	N/A
Default	Enable
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # no ib sm honor-partitions switch (config) # show ib sm honor-partitions disable switch (config) #</pre>
Related Commands	show ib sm honor-partitions
Notes	• If partitioning is disabled (no_partition_enforcement=1), then no named partitions can be enabled • If partitioning is enabled globally, the no_partition_enforcement changes from 1 to 0, and all defined partitions with state enabled are instantiated • If partitioning is globally disabled, all partitions are removed from the subnet, but the state (enabled or disabled) associated with defined partitions is not modified

ib sm hoq-lifetime

ib sm hoq-lifetime <time>

Sets the maximum time a frame can wait at the head of a switch-to-switch port queue before it is dropped.

Syntax Description	time	The time is 4.096 uS * 2time. The range of time is 0 to 20. A time of 20 means infinite, and the default value is 18 which translates to about 1 second.
Default	0x12 (~ 1 second)	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm hoq-lifetime 15 switch (config) # show ib sm hoq-lifetime 0xF (About 134 mS) switch (config) #</pre>	
Related Commands	show ib sm hoq-lifetime	
Notes		

ib sm ignore-other-sm

ib sm ignore-other-sm
no ib sm ignore-other-sm

Ignores all the rules governing SM elections and attempts to manage the fabric.

The no form of the command does not allow the SM to manage fabric if it loses the election.

Syntax Description	N/A
Default	Disable
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # ib sm ignore-other-sm switch (config) # show ib sm ignore-other-sm enable switch (config) #</pre>
Related Commands	show ib sm ignore-other-sm
Notes	

ib sm ipv6-nsm

ib sm ipv6-nsm
no ib sm ipv6-nsm

Consolidates IPv6 SNM group joins to 1 MC group per-MGID PKEY.
 The no form of the command disables the consolidation of IPv6 SNM.

Syntax Description	N/A
Default	Disable
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # ib sm ipv6-nsm switch (config) # show ib sm ipv6-nsm enable switch (config) #</pre>
Related Commands	show ib sm ipv6-nsm
Notes	

ib sm lash

ib sm lash {do-mesh-analysis | start-vl <vl-value>}
no ib sm lash do-mesh-analysis

Modifies “lash” routing method parameters.
The no form of the command disables SM “lash” routing for mesh analysis.

Syntax Description	do-mesh-analysis	Enables SM “lash” routing for mesh analysis.
	start-vl <vl-value>	Configures the starting VL for SM “lash” routing for mesh analysis (assuming that lash routing is enabled)
Default	do-mesh-analysis: disable start-vl: 0	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # ib sm lash do-mesh-analysis switch (config) # show ib sm lash do-mesh-analysis enable switch (config) #	
Related Commands	show ib sm lash do-mesh-analysis	
Notes		

ib sm leafhoq-lifetime

ib sm leafhoq-lifetime <time>

Sets the maximum time a frame can wait at the head of a switch-to-CA_or_Router port queue before it is dropped.

Syntax Description	time	The time is $4.096 \mu\text{S} * 2\text{time}$. The range of time is 0 to 20. A time of 20 means infinite, and the default value is 16 which translates to about 268 millisecond.
Default	0x10 (about 268 mS)	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm leafhoq-lifetime 8 switch (config) # show ib sm leafhoq-lifetime 0x8 (About 1 mS) switch (config) #</pre>	
Related Commands	show ib sm leafhoq-lifetime	
Notes		

ib sm leafvl-stalls

ib sm leafvl-stalls <count>

Sets the number of sequential frame drops that cause a switch-to-CA_or_Router port to enter the VLStalled state.

Syntax Description	count	1-255
Default	7	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm leafvl-stalls 3 switch (config) # show ib sm leafvl-stalls 3 switch (config) #</pre>	
Related Commands	show ib sm leafvl-stalls	
Notes		

ib sm lmc

ib sm lmc <mask>

Sets the LID Mask Control (LMC) value to be used on this subnet.

Syntax Description	mask	Valid values are 0-7.
Default	The default value is 0, which means that every port has exactly one unique LID.	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm lmc 7 switch (config) # show ib sm lmc 0x7 switch (config) #</pre>	
Related Commands	show ib sm lmc	
Notes		

ib sm lmc-esp0

ib sm lmc-esp0
no ib sm lmc-esp0

Sets the LMC for the subnet to be used for Enhanced Switch Port 0.

Syntax Description	N/A
Default	disable
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # ib sm lmc-esp0 switch (config) #
Related Commands	show ib sm lmc-esp0
Notes	

ib sm log-flags

ib sm log-flags [all] [debug] [error] [frames] [funcs] [info] [none] [routing] [verbose]
no ib sm log-flags

Controls what messages the SM logs.
 The no form of the command indicates to the SM not to run on this node.

Syntax Description	all	Turns on all the flags that follow (error info verbose debug funcs frames routing).
	debug	Logs diagnostic messages, high volume.
	error	Logs error messages.
	frames	Logs all SMP and GMP frames.
	funcs	Logs function entry/exit, very high volume.
	info	Logs basic messages, low volume.
	none	Turns off all logging flags.
	routing	Logs FDB routing information.
	verbose	Logs interesting stuff, moderate volume.
Default	0x3 (error, info)	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm log-flags error verbose funcs frames switch (config) # show ib sm log-flags 0x35 (error, verbose, funcs, frames) switch (config) #</pre>	
Related Commands	show ib sm log-flags	
Notes	• Every execution of this command replaces the current logging flags • The options “all” and “none” must be specified as the only parameter	

ib sm log-max-size

ib sm log-max-size <size>

Sets the maximum size of the log file to be <size> megabytes.

Syntax Description	size	Range: 1-60
Default	20 MBytes	
Configuration Mode	Config	
History	3.1.0000 3.5.1000	Updated Syntax Description, and Default
Role	admin	
Example	<pre>switch (config) # ib sm log-max-size 50 switch (config) # show ib sm log-max-size 50 MBytes switch (config) #</pre>	
Related Commands	show ib sm log-max-size	
Notes	- The log file “opensm_<switch_name>.log” is rotated when it exceeds the configured maximum file size up to 5 compressed files - When the log gets to the maximum size, or system storage fills up, the current log is deleted and messages start accumulating - To successfully upgrade from a version prior to 3.5.1000, this parameter must be set to a value in the range specified in the syntax description	

ib sm max-op-vls

ib sm max-op-vls <count>

Sets the maximum number of VLs supported on this subnet.

Syntax Description	count	Possible values: 1, 2, 4, 8, or 15.
Default	15	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # ib sm max-op-vls 4 switch (config) # show ib sm max-op-vls 4 switch (config) #	
Related Commands	show ib sm max-op-vls	
Notes		

ib sm max-reply-time

ib sm max-reply-time <time>

Sets the maximum time the SM waits for a reply before the transaction times out.

Syntax Description	time	Must be an integer (in milliseconds).
Default	200 milliseconds	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm max-reply-time 500 switch (config) # show ib sm max-reply-time 500 milliseconds switch (config) #</pre>	
Related Commands	show sm max-reply-time	
Notes		

ib sm max-reverse-hops

ib sm max-reverse-hops <max-reverse-hops>

Sets the maximum number of hops from the top switch to an I/O node.

Syntax Description	N/A
Default	0 hops
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # ib sm max-reverse-hops 500 switch (config) # show ib sm max-reverse-hops 500 hops switch (config) #
Related Commands	show ib sm max-reverse-hops
Notes	

ib sm max-wire-smps

ib sm max-wire-smps <count>

Sets the maximal number of MADs the SM has outstanding at one time to count.

Syntax Description	count	Number of concurrent mgmt packets. The value must be an integer.
Default	4	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm max-wire-smps 8 switch (config) # show ib sm max-wire-smps 8 switch (config) #</pre>	
Related Commands	show ib sm max-wire-smps	
Notes		

ib sm max-wire-smpps2

ib sm max-wire-smpps2 <count>

Sets the maximal timeout based outstanding SM management packets.

Syntax Description	count	Number of concurrent management packets. The value must be an integer.
Default	4	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm max-wire-smpps 8 switch (config) # show ib sm max-wire-smpps 8 switch (config) #</pre>	
Related Commands	show ib sm max-wire-smpps2	
Notes		

ib sm m-key

ib sm m-key <mkey>
no ib sm m-key

Configures the MKey used by the SM.
The no form of the command resets the MKey configuration to its default value.

Syntax Description	mkey	64-bit MKey
Default	00:00:00:00:00:00:00:00	
Configuration Mode	Config	
History	3.6.2002	
Role	admin	
Example	switch (config) # ib sm m-key 11:33:55:77:99:aa:cc:ee	
Related Commands	ib sm mkey-lease ib sm mkey-lookup ib sm mkey-protect-level show ib sm mkey-lease	
Notes	All nodes in the subnet may have to be reset or power-cycled after altering the SM MKey configuration	

ib sm mkey-lease

ib sm mkey-lease <time>
no ib sm mkey-lease

Configures the lease period used when MKey is non-zero.
The no form of the command resets this value to its default.

Syntax Description	time	MKey lease period in seconds Range: 0-65535; 0=unlimited
Default	0	
Configuration Mode	Config	
History	3.6.2002	
Role	admin	
Example	switch (config) # ib sm mkey-lease 660	
Related Commands	show ib sm mkey-lease	
Notes		

ib sm mkey-lookup

ib sm mkey-lookup
no ib sm mkey-lookup

Enables using a file cache (guid2mkey) to resolve unknown node MKey.
The no form of the command disables using a file cache to resolve unknown node MKey and the configured MKey is used for all ports.

Syntax Description	N/A
Default	Enabled
Configuration Mode	Config
History	3.6.2002
Role	admin
Example	switch (config) # ib sm mkey-lookup
Related Commands	show ib sm mkey-lookup
Notes	MKey lookup is a boolean value that controls how the SM finds the MKey of ports.

ib sm mkey-protect-level

ib sm mkey-protect-level <level>
no ib sm mkey-protect-level

Controls what data is returned to a get_PortInfo MAD request when the MKey in the request does not match the MKey on the port.
 The no form of the command resets the parameter to its default value.

Syntax Description	level	0 – when PortInfo is “read”, the actual MKey is returned in port info data 1 – when PortInfo is “read”, and the MKey in the MAD does not match the MKey on the port, the MKey value in the returned PortInfo data is set to 0. 2 – when PortInfo is “read”, and the MKey in the MAD does not match the MKey on the port, no data is returned.
Default	0	
Configuration Mode	Config	
History	3.6.2002	
Role	admin	
Example	switch (config) # ib sm mkey-protect-level 0	
Related Commands	show ib sm mkey-protect-level	
Notes		

ib sm msgfifo-timeout

ib sm msgfifo-timeout <time>

Sets the time value to be used by the subnet administrator to control when a BUSY status is returned to a client.

Syntax Description	time	In milliseconds.
Default	10 seconds	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm msgfifo-timeout 50000 switch (config) # show ib sm msgfifo-timeout 50.000 seconds switch (config) #</pre>	
Related Commands	show ib sm msgfifo-timeout	
Notes	If there is more than one message in the SA queue, and it has been there longer than time milliseconds, all additional incoming requests are immediately replied to with BUSY status.	

ib sm multicast

ib sm multicast
no ib sm multicast

Enables the SM to support multicasts on the fabric.
The no form of the command disables the SM from supporting multicasts on the fabric.

Syntax Description	N/A
Default	Disable
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # ib sm multicast switch (config) # show ib sm multicast enable switch (config) #
Related Commands	show ib sm multicast
Notes	

ib sm no-client-rereg

ib sm no-client-rereg
no ib sm no-client-rereg

Enables client re-registration requests.
The no form of the command disables client re-registration requests.

Syntax Description	N/A
Default	disable
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # ib sm no-client-rereg switch (config) # show ib sm no-client-rereg enable switch (config) #</pre>
Related Commands	show ib sm no-client-rereg
Notes	

ib sm overrun-trigger

ib sm overrun-trigger <count>

Enables SMA to generate standard InfiniBand trap number 130 when the number of local buffer overrun errors equals the count value, and the port's SMA supports traps.

Syntax Description	count	Range: 0-255.
Default	8	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm overrun-trigger 3 switch (config) # show ib sm overrun-trigger 3 switch (config) #</pre>	
Related Commands	show ib sm overrun-trigger	
Notes	Refer to the <i>InfiniBand Architecture Specification V1 r1.2.1</i> , section 14.2.5.1 table 131: Traps.	

ib sm packet-life-time

ib sm packet-life-time <time>

Sets the maximum time a frame can live in a switch.

Syntax Description	time	The time is 4.096 uS * 2*<time>. The rang is: 0-20. A time of 20 means infinite. The value 0x14 disables this mechanism.
Default	0x12 (about 1 second)	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm packet-life-time 20 switch (config) # show ib sm packet-life-time 0x14 (Infinite) switch (config) #</pre>	
Related Commands	show ib sm packet-life-time	
Notes		

ib sm phy-err-trigger

ib sm phy-err-trigger <count>

Enables SMA to generate trap 129 when the number of local link integrity errors equals the <count> value, and the port's SMA supports traps.

Syntax Description	count	Range is: 0-255.
Default	8	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # ib sm phy-err-trigger 5 switch (config) # show ib sm phy-err-trigger 5 switch (config) #	
Related Commands	show ib sm phy-err-trigger	
Notes		

ib sm polling-retries

ib sm polling-retries <value>

This variable defines the number of consecutive times an active SM must fail to respond before it is declared dead.

Syntax Description	value	Must be an integer.
Default	4	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm polling-retries 8 switch (config) # show ib sm polling-retries 8 switch (config) #</pre>	
Related Commands	show ib sm polling-retries	
Notes	The time between when the active SM fails and the time this SM declares it dead is: (sm_sminfo_polling_timeout * value) milliseconds.	

ib sm port-prof-switch

ib sm port-prof-switch
no ib sm port-prof-switch

Enables the counting of adapters, routers, and switches routed through links. The no form of the command disables the counting of adapters, routers, and switches routed through links.

Syntax Description	N/A
Default	False
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # ib sm port-prof-switch switch (config) # show ib sm port-prof-switch true switch (config) #
Related Commands	show ib sm port-prof-switch
Notes	

ib sm reassign-lids

ib sm reassign-lids
no ib sm reassign-lids

Controls the ability of the SM to reassign LIDs to nodes it finds already configured with a valid LID.

The no form of the command disables the SM from reassigning LIDs to nodes it finds already configured with a valid LID.

Syntax Description	N/A
Default	disable
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # ib sm reassign-lids switch (config) # show ib sm reassign-lids enable switch (config) #</pre>
Related Commands	show ib sm reassign-lids
Notes	• If enabled (ib sm reassign-lids), the SM can, but is not required to, reassign the LID on a node with a pre-configured LID • If disabled (no ib sm reassign-lids), the SM does not reassign LIDs • There are times when the SM is required to reassign LIDs or the fabric cannot be brought to a stable state, or a fabric option (like LMC) can not be fully applied

ib sm reset-config

ib sm reset-config

Resets all SM configuration options to defaults.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # ib sm reset-config switch (config) #
Related Commands	
Notes	

ib sm root-guid

ib sm root-guid <guid>
no ib sm root-guid <guid>

Adds a root GUID for the SM.
The no form of the command removes the GUID from the root GUID list.

Syntax Description	guid	The root GUID number in hexadecimal notation
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config)# ib sm root-guid aa:bb:00:11:22:33:44:55 switch (config) #	
Related Commands	show ib sm routing-engines	
Notes	The list of root GIDs are relevant when IB SM is running on the switch, and the routing algorithm is up-down or fat-tree.	

ib sm routing-engines

ib sm routing-engines [dor] [file] [ftree] [lash] [minhop] [none] [updn]
no ib sm routing-engines

Sets the routing engine of the SM.
 The no form of the command sets the routing engine to be “none”. The default SM routing engine is used.

Syntax Description	dor	Includes “dor” engine in selection of routing engines
	file	Includes “file” engine in selection of routing engines
	ftree	Includes “ftree” engine in selection of routing engines
	lash	Includes “lash” engine in selection of routing engines
	minhop	Includes “minhop” engine in selection of routing engines
	none	No routing engines specified; use SM default(s)
	updn	Includes “up/down” engine in selection of routing engines
Default	None	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # ib sm routing-engines none switch (config) # show ib sm routing-engines none none	
Related Commands	show ib sm routing-engines	
Notes	Multiple routing engines can be specified separated by# commas so that specific ordering of routing algorithms will be tried if earlier routing engines fail.	

ib sm rtr-aguid-enable

ib sm rtr-aguid-enable <value>
sm ib sm rtr-aguid-enable <value>

Configures SM alias GUID control option.
The no form of the command resets SM alias GUID control to its default value.

Syntax Description	value	Possible values: • 0 – does not configure alias GIDs required by routers • 1 – configures alias GIDs required by routers • 2 – clears and does not configure alias GIDs required by routers
Default	0	
Configuration Mode	Config	
History	3.6.2002	
Role	admin	
Example	switch (config) # ib sm rtr-aguid-enable 1	
Related Commands		
Notes		

ib sm rtr-pr-flow-label

ib sm rtr-pr-flow-label <value>
no ib sm rtr-pr-flow-label <value>

Configures inter-subnet PathRecord FlowLabel.
The no form of the command resets inter-subnet PathRecord FlowLabel to its default value.

Syntax Description	value	Range: 0-1048575
Default	0	
Configuration Mode	Config	
History	3.6.2002	
Role	admin	
Example	switch (config) # ib sm rtr-pr-flow-label 1	
Related Commands		
Notes		

ib sm rtr-pr-mtu

ib sm rtr-pr-mtu <value>
no ib sm rtr-pr-mtu <value>

Configures inter-subnet PathRecord MTU.
The no form of the command resets inter-subnet PathRecord MTU to its default value.

Syntax Description	value	Possible values: 256, 512, 1K, 2K, 4K
Default	2K	
Configuration Mode	Config	
History	3.6.2002	
Role	admin	
Example	switch (config) # ib sm rtr-pr-mtu 2k	
Related Commands		
Notes		

ib sm rtr-pr-rate

ib sm rtr-pr-rate <value>
no ib sm rtr-pr-rate <value>

Configures inter-subnet PathRecord rate.
The no form of the command resets inter-subnet PathRecord rate to its default value.

Syntax Description	value	Possible values: 2.5, 5, 10, 14, 20, 25, 40, 56, 100
Default	100	
Configuration Mode	Config	
History	3.6.2002	
Role	admin	
Example	switch (config) # ib sm rtr-pr-rate 5	
Related Commands		
Notes		

ib sm rtr-pr-sl

ib sm rtr-pr-sl <value>
no ib sm rtr-pr-sl <value>

Configures inter-subnet PathRecord SL.
The no form of the command resets inter-subnet PathRecord SL to its default value.

Syntax Description	value	Range: [0-15]
Default	0	
Configuration Mode	Config	
History	3.6.2002	
Role	admin	
Example	switch (config) # rtr-pr-sl 0	
Related Commands		
Notes		

ib sm rtr-pr-tclass

ib sm rtr-pr-tclass <value>
no ib sm rtr-pr-tclass <value>

Configures inter-subnet PathRecord T-class.
 The no form of the command resets inter-subnet PathRecord T-class to its default value.

Syntax Description	value	Range: [0-255]
Default	0	
Configuration Mode	Config	
History	3.6.2002	
Role	admin	
Example	switch (config) # ib sm rtr-pr-tclass 1	
Related Commands		
Notes		

ib sm sa-key

ib sm sa-key <SA_Key>

Sets the SA_Key 64-bit value used by SA to qualify that a query is “trusted”.

Syntax Description	SA Key	64 bit
Default	00:00:00:00:00:00:00:01	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm sa-key 5 switch (config) # show ib sm sa-key 00:00:00:00:00:00:00:05 switch (config) #</pre>	
Related Commands	show ib sm sa-key	
Notes	OpenSM version 3.2.1 and lower used the default value of “1” in host byte order. You may need to change this value to inter-operate with older subnet managers.	

ib sm single-thread

ib sm single-thread
no ib sm single-thread

Enables the Subnet Manager to use a single thread to service all requests.
The no form of the command enables SA to use multiple service threads.

Syntax Description	N/A
Default	Disable (use multiple service threads).
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # ib sm single-thread switch (config) # show ib sm single-thread enable switch (config) #
Related Commands	show ib sm single-thread
Notes	

ib sm sm-inactive

ib sm sm-inactive
no ib sm sm-inactive

Configures the SM to start in the “inactive” SM state. This option can be used to run a standalone system without the SM/SA function.
The no form of the command configures the SM to start in “init” SM state.

Syntax Description	N/A
Default	Disable
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # ib sm sm-inactive switch (config) # show ib sm sm-inactive enable switch (config) #
Related Commands	show ib sm sm-inactive
Notes	

ib sm sm-key

ib sm sm-key <SM_Key>

Sets the SM 64-bit SM_Key.

Syntax Description	SM Key	64 bit
Default	00:00:00:00:00:00:00:01	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm sm-key 00:00:00:00:00:00:00:05 switch (config) # show ib sm sm-key 00:00:00:00:00:00:00:05 switch (config) #</pre>	
Related Commands	show ib sm sm-key	
Notes	OpenSM version 3.2.1 and lower used the default value of “1” in host byte order. You may need to change this value to inter-operate with older subnet managers.	

ib sm sm-priority

ib sm sm-priority <priority>

Prioritizes the desired SM compared to other SMs on the fabric.

Syntax Description	priority	Priority 0 is the least important, 15 the most important.
Default	0	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm sm-priority 1 switch (config) # show ib sm sm-priority 1 switch (config) #</pre>	
Related Commands	show ib sm sm-priority	
Notes	If two or more active SMs have the same highest priority, the one with the lowest port GUID manages the fabric.	

ib sm sm-sl

ib sm sm-sl <sm-sl>

Sets the SM service level for SM/SA communication.

Syntax Description	sm-sl	0-15.
Default	0	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm sm-sl 10 switch (config) # show ib sm sm-sl 10 switch (config) #</pre>	
Related Commands	show ib sm sm-sl	
Notes	Selects the SL that is used for MADs.	

ib sm sminfo-poll-time

ib sm sminfo-poll-time <time>

This variable controls the timeout between two polls of an active subnet manager.

Syntax Description	time	In milliseconds.
Default	10 seconds	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # ib sm sminfo-poll-time 15 switch (config) # show ib sm sminfo-poll-time 15 milliseconds switch (config) #	
Related Commands	show ib sm sminfo-poll-time	
Notes		

ib sm subnet-prefix

ib sm subnet-prefix <prefix>
no ib sm subnet-prefix <prefix>

Sets the SM “Subnet Prefix” used to create scope qualifiers for all elements managed by the SM.

The no form of the command resets the subnet prefix to its default value.

Syntax Description	prefix	64 bit
Default	FE:80:00:00:00:00:00:00	
Configuration Mode	Config	
History	3.6.1002	
	3.6.2002	Added no form of the command
Role	admin	
Example	<pre>switch (config) # ib sm subnet-prefix ff:ff:ff:ff:ff:ff:ff:00 switch (config) # show ib sm subnet-prefix FF:FF:FF:FF:FF:FF:FF:00 switch (config) #</pre>	
Related Commands	show ib sm subnet-prefix	
Notes	The default value is also the InfiniBand default for a locally administered subnet.	

ib sm subnet-prefix-override

ib sm subnet-prefix-override
no ib sm subnet-override

Disables IB Router subnet prefix checking.
The no form of the command enables IB Router subnet prefix checking.

Syntax Description	N/A
Default	Enabled
Configuration Mode	Config
History	3.6.2002
Role	admin
Example	switch (config) # ib sm subnet-prefix-override switch (config) #
Related Commands	show ib sm subnet-prefix-override
Notes	

ib sm subnet-timeout

ib sm subnet-timeout <time>

Sets the global per-port subnet timeout value (PortInfo:SubnetTimeOut). This value also controls the maximum trap frequency in which no traps are allowed to be sent faster than the subnet_timeout value.

Syntax Description	time	The actual timeout is $4.096 \text{ uS} * 2^{*}\langle\text{time}\rangle$. The range of time is 0-31 for this parameter which supports 32 discrete time values between 4 uS and about 2.4 hours.
Default	0x12 (About 1 second)	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm subnet-timeout 5 switch (config) # show ib sm subnet-timeout 0x5 (About 131 uS) switch (config) #</pre>	
Related Commands	show ib sm subnet-timeout	
Notes	If the SMA generates a sequence of traps, the interval between successive traps should not be smaller than <time>.	

ib sm sweep-interval

ib sm sweep-interval <time>
no ib sm sweep-interval

Specifies the time between subnet sweeps.
The no form of the command disables periodic sweeps.

Syntax Description	time	Range: Between 0 and 36000 seconds (0 - disable).
Default	10 seconds	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # ib sm sweep-interval 20 switch (config) # show ib sm sweep-interval 20 seconds switch (config) #	
Related Commands	show ib sm sweep-interval	
Notes		

ib sm sweep-on-trap

ib sm sweep-on-trap
no ib sm sweep-on-trap

Enables every TRAP received by the SM to initiate a heavy sweep in addition to the processing required by the TRAP.

The no form of the command enables SM to use a combination of light and heavy sweeps based on the type of TRAP and other internal states.

Syntax Description	N/A
Default	enable
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # ib sm sweep-on-trap switch (config) # show ib sm sweep-on-trap enable switch (config) #</pre>
Related Commands	show ib sm sweep-on-trap
Notes	More than 10 successive identical TRAPs disable the automatic sweep behavior until at least one different TRAP has been received.

ib sm transaction-retries

ib sm transaction-retries <transaction-retries-count>

Sets the maximum retries for failed transactions.

Syntax Description	transaction-retries-count	Must be an integer.
Default	3	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm transaction-retries 10 switch (config) # show ib sm transaction-retries 10 switch (config) #</pre>	
Related Commands	show ib sm transaction-retries	
Notes		

ib sm use-heavy-sweeps

ib sm use-heavy-sweeps
no ib sm use-heavy-sweeps

Turns every fabric sweep to a heavy sweep.
The no form of the command enables the SM to use a combination of light and heavy sweeps.

Syntax Description	N/A
Default	disable
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # ib sm use-heavy-sweeps switch (config) # show ib sm use-heavy-sweeps enable switch (config) #
Related Commands	show ib sm use-heavy-sweeps
Notes	

ib sm use-ucast-cache

ib sm use-ucast-cache
no ib sm use-ucast-cache

Enables the SM to use cached routine data (LMC=0 only).
The no form of the command disables the SM to use cached routine data.

Syntax Description	N/A
Default	Disable
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # ib sm use-ucast-cache switch (config) # show ib sm use-ucast-cache true switch (config) #
Related Commands	show ib sm use-ucast-cache
Notes	

ib sm vl-stalls

ib sm vl-stalls <count>

Sets the number of sequential frame drops that cause a switch-to-switch port to enter the VLStalled state.

Syntax Description	count	1-255
Default	7	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib sm vl-stalls 10 switch (config) # show ib sm vl-stalls 10 switch (config) #</pre>	
Related Commands	show ib sm vl-stalls	
Notes		

ib sm virt

ib sm virt {enable | disable | ignore}
no ib sm virt

Configures IB SM port virtualization support.
The no form of the command resets this parameter to its default value.

Syntax Description	enable	IB SM supports virtualization, and configures virtual ports
	disable	IB SM disables virtual ports
	ignore	IB SM ignores virtual ports and does not change their configuration
Default	Ignore	
Configuration Mode	Config	
History	3.4.2008	
Role	admin	
Example	switch (config) # ib sm virt configure switch (config) #	
Related Commands		
Notes		

ib sm virt-default-hop-limit

ib sm virt-default-hop-limit <value>
no ib sm virt-default-hop-limit

Configures the default value for hop limit to be returned in path records.
The no form of the command resets this parameter to its default value.

Syntax Description	value	Range: 0-255
Default	2	
Configuration Mode	Config	
History	3.6.2002	
Role	admin	
Example	switch (config) # ib sm virt-default-hop-limit 3	
Related Commands		
Notes		

ib sm virt-max-ports-in-process

ib sm virt-max-ports-in-process <value>
no ib sm virt-max-ports-in-process

Configures the maximum number of ports to be processed simultaneously.
The no form of the command resets this parameter to its default value.

Syntax Description	value	Range:0-65535 '0' processes all pending ports
Default	4	
Configuration Mode	Config	
History	3.6.2002	
Role	admin	
Example	switch (config) # ib sm virt-max-ports-in-process 5	
Related Commands		
Notes		

show ib sm

show ib sm

Displays the SM admin state.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm enable switch (config) #
Related Commands	ib sm
Notes	

show ib sm accum-log-file

show ib sm accum-log-file

Displays the accum-log-file configuration.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm accum-log-file enable switch (config) #
Related Commands	ib sm accum-log-file
Notes	

show ib sm babbling-policy

show ib sm babbling-policy

Displays the ability of the SM to disable babbling ports (i.e., generating frequent traps).

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm babbling-policy disable switch (config) #
Related Commands	ib sm babbling-policy
Notes	

show ib sm connect-roots

show ib sm connect-roots

Displays the IBA compliant multi-stage switch directive.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # show ib sm connect-roots true switch (config) #</pre>
Related Commands	ib sm connect-roots
Notes	

show ib sm enable-quirks

show ib sm enable-quirks

Displays if the SM uses high risk features and handles HW workarounds.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm enable-quirks disable switch (config) #
Related Commands	ib sm enable-quirks
Notes	

show ib sm exit-on-fatal

show ib sm exit-on-fatal

Displays if the SM exits upon a fatal error.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm exit-on-fatal enable switch (config) #
Related Commands	ib sm exit-on-fatal
Notes	

show ib sm fdr10

show ib sm fdr10

Displays the status of the SM use of FDR10.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm fdr10 SM use of fdr10 is off switch (config) #
Related Commands	
Notes	

show ib sm force-link-speed

show ib sm force-link-speed

Displays SM behavior for PortInfo:LinkSpeedEnabled parameter on switch ports.

Syntax Description	N/A	
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
	3.4.1604	Updated Syntax Description, Example and Notess
Role	admin	
Example	<pre>switch (config) # show ib sm force-link-speed Default: set to PortInfo:LinkSpeedSupported switch (config) #</pre>	
Related Commands	ib sm force-link-speed	
Notes	Possible outputs: • Default: set to PortInfo:LinkSpeedExtSupported • Disabled: extended link speed not in use • Negotiate: <a list containing fdr, edr speeds>	

show ib sm force-link-speed-ext

show ib sm force-link-speed-ext

Displays SM behavior for PortInfo:LinkSpeedExtEnabled parameter on the switch ports.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000 3.4.1604 Updated description and Example
Role	admin
Example	switch (config) # show ib sm force-link-speed-ext Negotiate: fdr edr switch (config) #
Related Commands	
Notes	Possible outputs: • Default: set to PortInfo:LinkSpeedExtSupported • Disabled: extended link speed not in use • Negotiate: <a list containing fdr, edr speeds>

show ib sm force-log-flush

show ib sm force-log-flush

Displays if every log message generated forces the log to be flushed.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm force-log-flush enable switch (config) #
Related Commands	ib sm force-log-flush
Notes	

show ib sm guid2lid-cache

show ib sm guid2lid-cache

Displays whether or not the SM honors the cached GUID-to-LID mapping information.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm guid2lid-cache disable switch (config) #
Related Commands	ib sm guid2-lid-cache
Notes	

show ib sm honor-partitions

show ib sm honor-partitions

Displays the partition enforcement settings in the subnet.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # show ib sm honor-partitions disable switch (config) #</pre>
Related Commands	ib sm honor-partitions
Notes	

show ib sm hoq-lifetime

show ib sm hoq-lifetime

Displays the maximum time a frame can wait at the head of a switch-to-switch port queue before it is dropped.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm hoq-lifetime 0x12 (About 1 second) switch (config) #
Related Commands	ib sm hoq-lifetime
Notes	

show ib sm ignore-other-sm

show ib sm ignore-other-sm

Displays if the rules governing SM elections and attempt to manage the fabric on the node are ignored by the SM.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm ignore-other-sm enable switch (config) #
Related Commands	ib sm ignore-other-sm
Notes	

show ib sm ipv6-nsm

show ib sm ipv6-nsm

Displays the consolidation of IPv6 Solicited Node Multicast (SNM) group join requests.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm ipv6-nsm enable switch (config) #
Related Commands	ib sm ipv6-nsm
Notes	

show ib sm lash

show ib sm lash {do-mesh-analysis | start-vl}

Display 'lash' routing method parameters.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm lash do-mesh-analysis enable switch (config) #
Related Commands	ib sm lash
Notes	

show ib sm leafhoq-lifetime

show ib sm leafhoq-lifetime

Displays the maximum time a frame can wait at the head of a switch-to-CA_or_Router port queue before it is dropped.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm leafhoq-lifetime 0x10 (About 268 mS) switch (config) #
Related Commands	ib sm leafhoq-lifetime
Notes	

show ib sm leafvl-stalls

show ib sm leafvl-stalls

Displays the number of sequential frame drops that case a switch-to-CA_or_Router port to enter the VLStalled state.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm leafvl-stalls 7 switch (config) #
Related Commands	ib sm leafvl-stalls
Notes	

show ib sm lmc

show ib sm lmc

Displays the LID Mask Control (LMC) value to be used on this subnet.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm lmc 0x0 switch (config) #
Related Commands	ib sm lmc
Notes	

show ib sm lmc-esp0

show ib sm lmc-esp0

Displays whether the LMC for the subnet is also used for Enhanced Switch Port 0 (ib sm lmc-esp0) or if the LMC for ESP0 ports is 0 (no ib sm lmc-esp0).

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm lmc-esp0 enable switch (config) #
Related Commands	ib sm lmc-esp0
Notes	

show ib sm log

show ib sm log [continuous] [[not] [matching <reg-expression>]]

Displays IB SM event logs.

Syntax Description	continuous	Displays IB SM new event log messages as they arrive
	not	Displays IB SM new event logs that do not match a given regular expression.
	matching <regular expression>	Displays IB SM event log messages that match a given regular expression.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # show ib sm log Jul 18 12:00:40 165863 [48026660] 0x03 -> OpenSM 3.3.13.MLNX- _20121224_9b362db Jul 18 12:00:40 168685 [48026660] 0x80 -> OpenSM 3.3.13.MLNX- _20121224_9b362db Jul 18 12:00:40 170789 [48026660] 0x02 -> osm_vendor_init: 1000 pending umads specified Jul 18 12:00:40 175696 [48026660] 0x80 -> Entering DISCOVERING state Jul 18 12:00:40 249448 [48026660] 0x02 -> osm_vendor_bind: Binding to port 0x2c903008b0440 Jul 18 12:00:40 293959 [48026660] 0x02 -> osm_vendor_bind: Binding to port 0x2c903008b0440 Jul 18 12:00:40 296921 [48026660] 0x02 -> osm_vendor_bind: Binding to port 0x2c903008b0440 Jul 18 12:00:40 304702 [48026660] 0x02 -> osm_opensm_bind: Setting IS_SM on port 0x0002c903008b0440 Jul 18 12:00:40 399744 [4A85D4B0] 0x80 -> Entering MASTER state switch (config) #</pre>	
Related Commands	show ib sm log-flags	
Notes		

show ib sm log-flags

show ib sm log-flags

Displays what type of messages the SM will log.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm log-flags 0x3 (error, info) switch (config) #
Related Commands	ib sm log-flags
Notes	

show ib sm log-max-size

show ib sm log-max-size

Displays the maximum size of the log file.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm log-max-size 50 MBytes switch (config) #
Related Commands	is sm log-max-size
Notes	

show ib sm m-key

show ib sm m-key

Displays MKey value.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm m-key 50 MBytes switch (config) #
Related Commands	ib sm m-key
Notes	

show ib sm max-op-vls

show ib sm max-op-vls

Displays the maximum number of VLs supported on this subnet.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm max-op-vls 15 switch (config) #
Related Commands	ib sm max-op-vls
Notes	

show ib sm max-ports

show ib sm max-ports

Displays the number of CA ports SM can manage

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm max-ports 2048 switch (config) #
Related Commands	ib sm max-ports
Notes	

show ib sm max-reply-time

show ib sm max-reply-time

Displays the maximum time in milliseconds that the SM will wait for a reply before the transaction times out.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm max-reply-time 200 milliseconds switch (config) #
Related Commands	ib sm max-reply-time
Notes	

show ib sm max-reverse-hops

show ib sm max-reverse-hops

Displays max hops IO node to top switch

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm max-reverse-hops 0 hops switch (config) #
Related Commands	ib sm max-reverse-hops
Notes	

show ib sm max-smps-timeout

show ib sm max-smps-timeout

Displays timeout for SMPs between max_wire_smps & max_wire_smps2.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm max-smps-timeout 600000 switch (config) #
Related Commands	
Notes	

show ib sm max-wire-smpls

show ib sm max-wire-smpls

Displays the maximal number of MADs the SM will have outstanding at one time to count.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm max-wire-smpls 8 switch (config) #
Related Commands	ib sm max-wire-smpls
Notes	

show ib sm max-wire-smpps2

show ib sm max-wire-smpps2

Displays maximal SM timeout based packets allowed to be outstanding.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm max-wire-smpps2 4 switch (config) #
Related Commands	
Notes	

show ib sm mkey-lease

show ib sm mkey-lease

Displays MKey period in seconds.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm mkey-lease 0 (No timeout) switch (config) #
Related Commands	ib sm mkey-lease
Notes	

show ib sm m-key

show ib sm m-key

Displays the MKey used by the SM

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.6.2002
Role	admin
Example	switch (config) # show ib sm m-key 11:33:55:77:99:aa:cc:ee
Related Commands	ib sm m-key
Notes	

show ib sm mkey-lookup

show ib sm mkey-lease

Displays whether SM looks in file cache for unknown node MKeys or not.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.6.2002
Role	admin
Example	switch (config) # show ib sm mkey-lookup enable
Related Commands	ib sm mkey-lookup
Notes	

show ib sm mkey-protect-level

show ib sm mkey-protect-level

Displays MKey protection level.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.6.2002
Role	admin
Example	switch (config) # show ib sm mkey-protect-level 0
Related Commands	ib sm mkey-protect-level
Notes	

show ib sm msgfifo-timeout

show ib sm msgfifo-timeout

Displays the elapsed time in milliseconds before a frame at the head of Subnet Agent queue causes an immediate BUSY state.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm msgfifo-timeout 10.000 seconds switch (config) #
Related Commands	ib sm msgfifo-timeout
Notes	

show ib sm multicast

show ib sm multicast

Displays whether the SM supports multicast on the fabric.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm multicast enable switch (config) #
Related Commands	ib sm multicast
Notes	

show ib sm no-client-rereg

show ib sm no-client-rereg

Displays client re-registration admin state.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm no-client-rereg enable switch (config) #
Related Commands	ib no-client-rereg
Notes	

show ib sm overrun-trigger

show ib sm overrun-trigger

Displays count of local buffer overrun errors for Infiniband trap 130

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm overrun-trigger 3 switch (config) #
Related Commands	ib sm overrun-trigger
Notes	

show ib sm packet-life-time

show ib sm packet-life-time

Displays the maximum time a frame can live in a switch.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm packet-life-time 0x14 (Infinite) switch (config) #
Related Commands	ib sm packet-life-time
Notes	

show ib sm phy-err-trigger

show ib sm phy-err-trigger

Displays the number of local link integrity errors and the port's SMA supports traps.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm phy-err-trigger 5 switch (config) #
Related Commands	ib sm phy-err-trigger
Notes	

show ib sm polling-retries

show ib sm polling-retries

Displays the number of consecutive times an active SM must fail to respond before it is declared dead.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm polling-retries 8 switch (config) #
Related Commands	ib sm polling-retries
Notes	

show ib sm port-prof-switch

show ib sm port-prof-switch

Displays whether or not the counting of adapters, routers, and switches through the links is being done.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm port-prof-switch true switch (config) #
Related Commands	ib sm port-prof-switch
Notes	

show ib sm reassign-lids

show ib sm reassign-lids

Displays the ability of the SM to reassign LIDs to nodes it finds already configured with a valid LID.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm reassign-lids enable switch (config) #
Related Commands	ib sm reassign-lids
Notes	

show ib sm root-guid

show ib sm root-guid

Displays the configured root GUIDs for the SM.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config)# show ib sm root-guid AA:00:11:22:33:44:55 AA:00:11:22:33:44:56 AA:00:11:22:33:44:57 ... switch (config)#</pre>
Related Commands	ib sm routing-engine
Notes	The list of root GUIDs are relevant when IB SM is running on the switch, and the routing algorithm is up-down or fat-tree.

show ib sm routing-engines

show ib sm routing-engines

Displays an ordered list of routing engines

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm routing-engine none switch (config) #
Related Commands	ib sm routing-engine
Notes	

show ib sm routing-info

show ib sm routing-info

Displays current routing engine information.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm routing-info Current routing engine minhop switch (config) #
Related Commands	
Notes	

show ib sm rtr-aguid-enable

show ib sm rtr-aguid-enable

Displays GUID option configuration.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.6.2002
Role	admin
Example	switch (config) # show ib sm rtr-aguid-enable 0
Related Commands	ib sm rtr-aguid-enable
Notes	

show ib sm rtr-pr-flow-label

show ib sm rtr-pr-flow-label

Displays inter-subnet PathRecord FlowLabel.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.6.2002
Role	admin
Example	switch (config) # show ib sm rtr-pr-flow-label 0
Related Commands	ib sm rtr-pr-flow-label
Notes	'0' means Inter-subnet PathRecord FlowLabel is disabled

show ib sm rtr-pr-mtu

show ib sm rtr-pr-mtu

Displays inter-subnet PathRecord MTU.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.6.2002
Role	admin
Example	switch (config) # show ib sm rtr-pr-mtu 2K
Related Commands	ib sm rtr-pr-mtu
Notes	

show ib sm rtr-pr-rate

show ib sm rtr-pr-rate

Displays inter-subnet PR rate.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.6.2002
Role	admin
Example	switch (config) # show ib sm rtr-pr-rate 100
Related Commands	ib sm rtr-pr-rate
Notes	

show ib sm rtr-pr-sl

show ib sm rtr-pr-sl

Displays inter-subnet PathRecord service level.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.6.2002
Role	admin
Example	switch (config) # show ib sm rtr-pr-sl 0
Related Commands	ib sm rtr-pr-sl
Notes	

show ib sm sa-key

show ib sm sa-key

Displays the SM sa-key value used by SA to qualify that a query is “trusted”.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm sa-key 00:00:00:00:00:00:00:05 switch (config) #
Related Commands	ib sm sa-key
Notes	

show ib sm single-thread

show ib sm single-thread

Displays if the SM uses a single thread to service all requests.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm single-thread enable switch (config) #
Related Commands	ib sm single-thread
Notes	

show ib sm sm-inactive

show ib sm sm-inactive

Displays whether or not the SM starts in “inactive” rather than “init” SM state.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm sm-inactive enable switch (config) #
Related Commands	ib sm sm-inactive
Notes	

show ib sm sm-key

show ib sm sm-key

Displays the SM 64-bit SM_Key.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm sm-key 00:00:00:00:00:00:00:05 switch (config) #
Related Commands	ib sm sm-key
Notes	

show ib sm sm-priority

show ib sm sm-priority

Displays the importance of this SM compared to other SMs on the fabric.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # show ib sm sm-priority 1 switch (config) #</pre>
Related Commands	ib sm sm-priority
Notes	Priority 0 is the least important, 15 the most important. If 2 or more active SMs have the same highest priority, the one with the lowest port GUID will manage the fabric.

show ib sm sm-sl

show ib sm sm-sl

Display SL used for SM/SA communication

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm sm-sl 1 switch (config) #
Related Commands	ib sm sm-sl
Notes	

show ib sm sminfo-poll-time

show ib sm sminfo-poll-time

Displays the timeout in milliseconds between two polls of an active SM.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm sminfo-poll-time 15 milliseconds switch (config) #
Related Commands	ib sm sminfo-poll-time
Notes	

show ib sm subnet-prefix

show ib sm subnet-prefix

Displays the SM “Subnet Prefix” used to create scope qualifiers for all elements managed by the SM.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm subnet-prefix FF:FF:FF:FF:FF:FF:FF:00 switch (config) #
Related Commands	ib sm subnet-prefix
Notes	

show ib sm subnet-prefix-override

show ib sm subnet-prefix

Displays whether IB Router subnet prefix checking is enabled or disabled.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.6.2002
Role	admin
Example	switch (config) # show ib sm subnet-prefix-override disable
Related Commands	ib sm subnet-prefix-override
Notes	

show ib sm subnet-timeout

show ib sm subnet-timeout

Displays the global per-port subnet timeout value (PortInfo:SubnetTimeOut). This value also controls the maximum trap frequency in which no traps are allowed to be sent faster than the subnet_timeout value. The time is 4.096 uS * 2*time.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm subnet-timeout 0x5 (About 131 uS) switch (config) #
Related Commands	ib sm subnet-timeout
Notes	

show ib sm sweep-interval

show ib sm sweep-interval

Displays the time in seconds between subnet sweeps.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm sweep-interval 20 seconds switch (config) #
Related Commands	ib sm sweep-interval
Notes	

show ib sm sweep-on-trap

show ib sm sweep-on-trap

Displays whether or not a heavy sweep is initiated by the TRAP received by the SM.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm sweep-on-trap enable switch (config) #
Related Commands	ib sm sweep-on-trap
Notes	

show ib sm transaction-retries

show ib sm transaction-retries

Displays maximum retries before failing a transaction

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm transaction-retries 3 switch (config) #
Related Commands	ib sm transaction-retries
Notes	

show ib sm use-heavy-sweeps

show ib sm use-heavy-sweeps

Displays SM requirement to always use heavy sweeps.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm use-heavy-sweeps disable switch (config) #
Related Commands	ib sm use-heavy-sweeps
Notes	

show ib sm use-ucast-cache

show ib sm use-ucast-cache

Displays if the SM uses cached routine data.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm use-ucast-cache false switch (config) #
Related Commands	ib sm user-ucase-cache
Notes	

show ib sm version

show ib sm version

Displays the open SM version that is currently running.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.2.3000
Role	admin
Example	switch (config) # show ib sm version OpenSM3.3.7 switch (config) #
Related Commands	
Notes	

show ib sm virt-default-hop-limit

show ib sm virt-default-hop-limit

Displays the default value for hop limit to be returned in path records.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.6.2002
Role	admin
Example	switch (config) # show ib sm virt-default-hop-limit 2
Related Commands	ib sm virt-default-hop-limit
Notes	

show ib sm virt-max-ports-in-process

show ib sm virt-max-ports-in-process

Displays the maximum number of ports to be processed simultaneously.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.6.2002
Role	admin
Example	switch (config) # show ib sm virt-max-ports-in-process 4
Related Commands	ib sm virt-max-ports-in-process
Notes	

show ib sm vl-stalls

show ib sm use-vl-stalls

Displays the number of sequential frame drops that cause a switch-to-switch port to enter the VLStalled state.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # show ib sm vl-stalls 7 switch (config) #
Related Commands	ib sm vl-stalls
Notes	

7.5.4.2 Partition (SM)

ib partition

ib partition <partition-name> [pkey <pkey number>]
no ib partition <partition-name>

Enters the context of the partition specified.
The no form of the command deletes the partition.

Syntax Description	partition-name	Name of partition context to be entered.
	pkey	Creates a partition and enters a new configuration mode.
Default	Default partition is available (PKEY 0x7fff)	
Configuration Mode	Config	
History	3.2.0500	
Role	admin	
Example	switch (config) # ib partition my-partition switch (config partition my-partition) #	
Related Commands		
Notes		

pkey

pkey <number> [force]
no pkey <number>

Specifies PKEY number for this partition.
 The no form of the command removes the PKEY configuration from partitions.conf file.

Syntax Description	number	0x001-0x7fff
	force	Forces configuration
Default	N/A	
Configuration Mode	Config Partition	
History	3.2.0500	
	3.5.1000	Added “force” parameter
Role	admin	
Example	<pre>switch (config) # ib partition my-partition switch (config partition my-partition) # pkey 0x7777 switch (config partition my-partition) #</pre>	
Related Commands		
Notes	PKEY must be unique.	

defmember

defmember <type> [force]
no defmember

Sets the default membership for port GUID list.
The no form of the command set the defmember configuration to default (it will not appear in the partitions.conf file).

Syntax Description	type	Default membership for GUIDs in this partition: • full • limited • both
	force	Forces configuration
Default	limited	
Configuration Mode	Config Partition	
History	3.2.0500	
	3.4.1100	Added “both” option
	3.5.1000	Added “force” parameter
Role	admin	
Example	switch (config) # ib partition my-partition switch (config ib partition my-partition) # defmember full switch (config ib partition my-partition) #	
Related Commands	ib sm allow-both-pkeys member	
Notes	This parameter can be overwritten for specific GUID, using the “member” command.	

member

member {<guid> | all | self} [type <member-type>] [force]
no member {<guid> | all | self} [type] [force]

Adds static members to partition.

The no form of the command will remove the static member from the partition (it will not appear in the partitions.conf file).

Syntax Description	guid	The GUID number
	all self	The options “all” can be used for all GUIDs in the fabric, or “self” for the switch guide
	member-type	Default membership for GUIDs in this partition: • full • limited • both
	force	Forces configuration
Default	N/A	
Configuration Mode	Config Partition	
History	3.2.0500	
	3.4.1100	Added “both” parameter
	3.5.1000	Added “force” parameter
Role	admin	
Example	<pre>switch (config) # ib partition my-partition switch (config partition my-partition) # member all switch (config partition my-partition) #</pre>	
Related Commands	<pre>ib sm allow-both-pkeys defmember</pre>	
Notes		

ipoib

ipoib [force]
no ipoib

Enables this partition to use IPoIB. As a result IPoIB multicast group will be created.

The no form of the command will remove the use of IPoIB in this partition (it will not appear in the partitions.conf file).

Syntax Description	force	Forces configuration
Default	no ipoib	
Configuration Mode	Config Partition	
History	3.2.0500	
	3.5.1000	Added “force” parameter
Role	admin	
Example	<pre>switch (config) # ib partition my-partition switch (config partition my-partition) # ipoib switch (config partition my-partition) #</pre>	
Related Commands	rate mtu sl scope	
Notes	“rate”, “mtu”, “sl” and “scope” commands can be used only when the IPoIB parameter is enabled.	

mtu

mtu <256, 512, 1K, 2K,4K> [force]
no mtu

Specifies MTU for this IPoIB multicast group.
 The no form of the command sets the mtu to default (it will not appear in the partitions.conf file).

Syntax Description	force	Forces configuration
Default	2K	
Configuration Mode	Config Partition	
History	3.2.0500	
	3.5.1000	Added “force” parameter
Role	admin	
Example	<pre>switch (config) # ib partition my-partition switch (config partition my-partition) # mtu 4K switch (config partition my-partition) #</pre>	
Related Commands	ipoib	
Notes	IPoIB parameter on the partitions must be enabled in order to use this parameter	

rate

rate <rate> [force]
no rate

Specifies rate for this IPoIB multicast group.
The no form of the command set the rate to default (removes the rate from the partitions.conf)

Syntax Description	rate	• default - Default • 2.5 - 2.5 Gbps • 5 - 5 Gbps • 10 - 10 Gbps • 14 - 14 Gbps • 20 - 20 Gbps • 25 - 25 Gbps • 40 - 40 Gbps • 56 - 56 Gbps • 100 - 100 Gbps
	force	Forces configuration
Default	10 Gbps.	
Configuration Mode	Config Partition	
History	3.2.0500	
	3.4.1100	Updated rate Syntax Description
	3.5.1000	Added “force” parameter
Role	admin	
Example	<pre>switch (config) # ib partition my-partition switch (config partition my-partition) # rate 20 switch (config partition my-partition) #</pre>	
Related Commands		
Notes	• Ports that do not support the IPoIB rate are not added to the partition	

scope

scope <type> [force]

no scope <link-local, site-local, organization-local, global>

Specifies scope for this IPoIB multicast group.

The no form of the command removes the scope configuration from the partitions.conf file

Syntax Description	type	link-local site-local organization-local global
	force	Forces configuration
Default	link-local	
Configuration Mode	Config Partition	
History	3.2.0500	
	3.5.1000	Added “force” parameter
Role	admin	
Example	<pre>switch (config) # ib partition my-partition switch (config partition my-partition) # scope global switch (config partition my-partition) #</pre>	
Related Commands		
Notes	ipoib parameter on the partitions must be enabled in order to use this parameter.	

sl

sl <0-14, “default”> [force]
no sl

Specifies SL (Service Level - QoS) for this IPoIB multicast group.
The no form of the command sets it to default (the sl configuration is removed from the partitions.conf file).

Syntax Description	0-14
	force Forces configuration
Default	default (0)
Configuration Mode	Config Partition
History	3.2.0500
	3.5.1000 Added “force” parameter
Role	admin
Example	switch (config) # ib partition my-partition switch (config partition my-partition) # sl 7 switch (config partition my-partition) #
Related Commands	
Notes	ipoib parameter on the partitions must be enabled in order to use this parameter.

show ib partition

show ib partition [<partition-name> [member [<member-name>]]]

Displays partition info, with optional to filters.

Syntax Description	partition-name	Filter the output per partition name.
	member <member-name>	Filter the output by a specific member
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.2.0500	
Role	admin	
Example	<pre>switch (config) # show ib partition Default PKey = 0x7FFF defmember = full ipoib = yes members GUID='ALL' member='full' my-partition PKey = Auto-assign defmember = full ipoib = yes rate = 40 mtu = 1K sl = 3 scope = link_local members GUID='ALL' member='default' GUID='SELF' member='full' switch (config) #</pre>	
Related Commands		
Notes		

7.5.4.3 Quality of Service (SM)

ib baseqos <port-type> high-limit

ib baseqos <port-type> high-limit <count>

Sets the high-limit value for the indicated port type. Thus the system will send at least $4096 * \text{<count>}$ bytes from the high priority list before sending any from the low priority list.

Syntax Description	port-type	- ca - channel adapters - rtr - routers - sw0 - ports 0 only of the switches - swe - external ports of the switches
	high-limit	Possible values are: -1...255 -1 - default SM high-limit 0 - 1 frame $i = 1 \dots 254 - 4K * i$ 255 - unlimited
Default	-1 (default SM high-limit).	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib baseqos ca high-limit 255 switch (config) #</pre>	
Related Commands	show ib baseqos	
Notes	A high-limit value of 255 means unlimited, and that makes it possible to starve the low priority list.	

ib baseqos <port-type> max-vls <value>

ib baseqos <port-type> max-vls <value>

Sets the maximum number of VLs for the indicated port type.

Syntax Description	port-type	ca - channel adapters rtr - routers sw0 - ports 0 only of the switches swe - external ports of the switches
	value	Max VLs range between 1 and 15.
Default	15	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib baseqos ca max-vls 15 switch (config) # show ib baseqos ca max-vls 15 switch (config) #</pre>	
Related Commands	show ib baseqos	
Notes		

ib baseqos <port-type> sl2vl

ib baseqos <port-type> sl2vl {sl0 | sl0 sl1 | sl0 sl1 sl2 |...}
no ib baseqos <port-type> sl2vl

Sets a list of up to 16 entries that map the SL entry to an appropriate VL.
The no form of the command sets the attributes to their default settings.

Syntax Description	port-type	ca - channel adapters rtr - routers sw0 - ports 0 only of the switches swe - external ports of the switches
	sl[i]	A single vector (1 ... 16 elements), the command line vector determine the SL [0...15] that is mapped to the specified VL [0...15].
Default	The default mapping is: 0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,7	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) #show ib baseqos ca sl2vl 0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,7 switch (config) # ib baseqos ca sl2vl 10 10 10 switch (config) # show ib baseqos ca sl2vl 10,10,10,15,15,15,15,15,15,15,15,15,15,15,15,15 switch (config) #</pre>	
Related Commands	show ib baseqos	
Notes	Any missing SLs will be mapped to VL15.	

ib baseqos <port-type> vlarb-high <value>

ib baseqos <port-type> vlarb-high {VW1 | VW1 VW2 | ...}
no ib baseqos <port-type> vlarb-high

Sets up to 15 VL to Weight mapping pairs for high priority processing.
 The no form of the command sets the attributes to their default settings.

Syntax Description	port-type	ca - channel adapters rtr - routers sw0 - ports 0 only of the switches swe - external ports of the switches
	VW[i]	There are two possible options for this parameter: - A single vector (1 ...15) in the format of “#:#” separated by spaces, see example below. - Format of “i#=X:Y” in order to change a specific entry (see example below)
Default	The default mapping is: 0:4,1:0,2:0,3:0,4:0,5:0,6:0,7:0,8:0,9:0,10:0,11:0,12:0,13:0,14:0	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) #show ib baseqos ca vlarb-high 0:4,1:0,2:0,3:0,4:0,5:0,6:0,7:0,8:0,9:0,10:0,11:0,12:0,13:0,14:0 switch (config) # ib baseqos ca vlarb-high 0:10 1:10 switch (config) # show ib baseqos ca vlarb-high 0:10,1:10,2:0,3:0,4:0,5:0,6:0,7:0,8:0,9:0,10:0,11:0,12:0,13:0,14:0 switch (config) # ib baseqos sw0 vlarb-high i2=4:3 switch (config) # show ib baseqos sw0 vlarb-high 0:10,1:10,4:3,3:0,4:0,5:0,6:0,7:0,8:0,9:0,10:0,11:0,12:0,13:0,14:0</pre>	
Related Commands	show ib baseqos	
Notes	- Unspecified elements will be filled with (index:0) - You may have multiple entries with the same VL on this list.	

ib baseqos <port-type> vlarb-low <value>

ib baseqos <port-type> vlarb-low {VW1 | VW1 VW2 | ...}
no ib baseqos <port-type> vlarb-low

Sets up to 15 VL to Weight mapping pairs for low priority processing.
The no form of the command sets the attributes to their default settings.

Syntax Description	port-type	ca - channel adapters rtr - routers sw0 - ports 0 only of the switches swe - external ports of the switches
	VW[i]	There are two possible options for this parameter: - A single vector (1 ...15) in the format of “#:#” separated by spaces, see example below. - Format of “i#=X:Y” in order to change a specific entry (see example below)
Default	The default mapping is: 0:0,1:4,2:4,3:4,4:4,5:4,6:4,7:4,8:4,9:4,10:4,11:4,12:4,13:4,14:4	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # ib baseqos sw0 vlarb-low 1:1 switch (config) # show ib baseqos sw0 vlarb-low 1:1, 1:0, 2:0, 3:0, 4:0, 5:0, 6:0, 7:0, 8:0, 9:0, 10:0, 11:0, 12:0, 13:0, 14:0 switch (config) # ib baseqos sw0 vlarb-low i2=4:3 switch (config) # show ib baseqos sw0 vlarb-low 1:1, 1:0, 4:3, 3:0, 4:0, 5:0, 6:0, 7:0, 8:0, 9:0, 10:0, 11:0, 12:0, 13:0, 14:0 switch (config) #</pre>	
Related Commands	show ib baseqos	
Notes	You may have multiple entries with the same VL on this list.	

ib baseqos reset-config

ib baseqos reset-config

Resets all basic QoS configuration options to defaults.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # ib baseqos reset-config switch (config) #
Related Commands	
Notes	

show ib baseqos

show ib baseqos <port-type> <baseqos-parameters>

Displays the base ib QoS configuration.

Syntax Description	port-type	- ca - channel adapters - rtr - routers - sw0 - ports 0 only of the switches - swe - external ports of the switches
	baseqos-parameters	Possible values are: - high-limit - Display high limit (how many high pri before low) - max-vls - Display maximum number of VLs supported on CAs in subnet - sl2vl - Display current SL-to-VL mapping vector - vlarb-high - Display current high priority VL arbitration - vlarb-low - Display current low priority VL arbitration
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # show ib baseqos ca high-limit 0 switch (config) #</pre>	
Related Commands		
Notes		

ib qos

ib qos
no ib qos

Enables advanced QoS management on this node
The no form of the command disables advance QoS on this node.

Syntax Description	N/A
Default	advance qos is disabled.
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	switch (config) # ib qos switch (config) # show ib qos enable switch (config) #
Related Commands	show ib qos
Notes	

ib qos level

ib qos level {<name> | default} {mtu-limit <mtu> | packet-life <time> | pkey <number> | rate-limit <rate-value> | sl <sl-value> | use <description>}
no ib qos level {<name> | default} {mtu-limit | packet-life | pkey | rate-limit | sl | use}

Specifies a QoS level <name> or “default” parameters.
The no form of the command set the parameters to default.

Syntax Description	<name> default	Specify a name for this qos group, or use the “default” for the default qos parameters.
	mtu-limit <mtu>	MTU in bytes. Possible values are: 1k, 256, 2k, 4k, 512
	packet-life <time>	Time a packet can wait in switch egress queue before being dropped. The bytes from 4 microsecond up to 2 seconds or infinite. Possible values are 0-20 0 - 4usec 1 - 8usec ... 20 - unlimited
	pkey <number>	PKEY value: ranges between -1 and 32767 (hex 0x7fff)
	rate-limit <rate-value>	Manages rate limits for QoS Policy levels. Possible values are (in Gbps): default, 2.5, 5, 10, 14, 20, 25, 40, 56, 100.
	sl <sl-value>	Manages service level for QoS Policy levels. Range: 0-15.
	use <description>	Specify usage description for this QoS level
Default	The default values are: use = “default QoS Level” sl = 0 mtu-limit = default rate-limit = default packet-life = 0x12 pkey = -1	
Configuration Mode	Config	
History	3.1.0000	
	3.4.1100	Updated description of “rate-limit” parameter

Role	admin
Example	<pre>switch (config) # ib qos level my-qos-group mtu-limit 2K switch (config) # show ib qos my-qos-group my-qos-group: use = default QoS Level sl = 0 mtu-limit = 2K rate-limit = default packet-life = 0x12 pkey = -1 switch (config) #</pre>
Related Commands	show ib qos
Notes	

ib qos match-rule

```

ib qos match-rule <rule-index> { {destination | source} <string> | {pkey | qos-
class | service-id} <index> {first | last} <value>} | qos-level-name <name>| use
<description>}
no ib qos match-rule <rule-index> { {destination | source} | {pkey | qos-
class | service-id} <index> {first | last} } | qos-level-name | use }

```

Manages QoS Policy match rules.

The no form of the command set the QoS match-rule to default.

Syntax Description	rule-index	Index of this match-rule. Possible range is: 0-4294967295
	destination source <string>	Manages destination or source for QoS Policy match rules.
	pkey qos-class ser- vice-id <index>	Manages values for QoS Policy match rules.
	{first last} <value>	First or last value range (per PKEY / qos-class of service id.
	qos-level-name <name>	Name for the QoS level
	use <description>	Specify usage description for this QoS level
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre> switch (config) # ib qos match-rule 10 use my-use switch (config) # show ib qos match-rule 10 match-rule/10: match-rules: use = my-use match-rules: qos-level-name = DEFAULT switch (config) # </pre>	
Related Commands	show ib qos	
Notes		

ib qos port-group

ib qos port-group <name> {node-type <index> type <node-type> | partition <name> | pkey <number> | port-guid <index> {first | last} <value> | port-name <index> name <name-value> | use <description>}
no ib qos port-group <name> {node-type <index> type | partition | pkey | port-guid <index> {first | last} | port-name <index> name | use }

Manages QoS Policy port groups.
 The no form of the command removes a QoS port-group.

Syntax Description	<name>	Port group name
	node-type <index>	Node type index
	type <node-type>	A node type for this port group
	partition <name>	A Partition name
	pkey <number>	A PKEY number
	port-guid <index> {first last} <value>	Port-guid range
	port-name <index> name <name-value>	Port index name
	use <description>	Specify usage description for this QoS level
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config)# ib qos port-group my-group use my-use switch (config)# show ib qos port-group my-group port-group/my-group: port-groups: pkey = -1 port-groups: use = my-use switch (config)#</pre>	
Related Commands	show ib qos	
Notes		

ib qos ulp any

ib qos ulp any {pkey | service-id | target-port-guid <index> {first | last | sl} <value> | sl <sl-value>}
no ib qos ulp any {pkey | service-id | target-port-guid <index> {first | last | sl} | sl}

Configures ULP any attributes.

The no form of the command deletes ULP any attributes.

Syntax Description	pkey <index>	Manages ULP default PKEY assignment.
	service-id <index>	Manages default ULP Service ID match rule.
	target-port-guid <index>	Manages ULP default target port GUID rule.
	first last sl <value>	- first - first value in range - last - last value n range - sl - Service level for the ULP rule
	sl <sl-value>	Sets default SL.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # ib qos ulp any sl 2 switch (config) #	
Related Commands	show ib qos	
Notes		

ib qos ulp ipoib

ib qos ulp ipoib {default sl <sl-value>| pkey <index> {first | last | sl} <value> }
no ib qos ulp ipoib {default sl | pkey <index>}

Manages ULP IPoIB settings.

The no form of the command deletes IPoIB settings.

Syntax Description	default sl <sl-value>	Set the default sl. Range 1-15
	pkey <index>	Manages ULP default PKEY assignment.
	first last sl <value>	- first - first value in range - last - last value n range - sl - Service level for the ULP rule
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Role	admin	
Example	switch (config) # ib qos ulp ipoib default sl 5 switch (config) #	
Related Commands	show ib qos	
Notes		

ib qos ulp <protocol-type>

ib qos ulp <protocol-type> {default sl <sl-value> | port-num< index> <first | last | sl> <value>}

no ib qos ulp iser {default <sl> | port-num1 <first | last | sl>}

Configures ULP IScsi Extensions for RDMA, Reliable Datagram Sockets or Sockets Direct Protocol attributes.

The no form of the command deletes all rules.

Syntax Description	protocol-type	iser - Scsi Extensions for RDMA rds - Reliable Datagram Sockets sdp - Sockets Direct Protocol
	default sl <sl-value>	Set the default sl. Range 1-15
	port-num< index>	Port number index
	first last sl	- first - First in range - last - in range - sl - Service level for the ULP rule
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # ib qos ulp iser default sl 2 switch (config) #	
Related Commands	show ib qos	
Notes		

ib qos ulp srp

ib qos ulp srp target-port-guid <index> <first | last | sl> <value>
no ib qos ulp srp target-port-guid <index>

Configures Scsi Rdma Protocol attributes
 The no form of the command deletes the rules.

Syntax Description	target-port-guid <index>	the index of the target port guid.
	first last sl	• first - First in range • last - in range • sl - Service level for the ULP rule
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # ib qos ulp srp target-port-guid 1 sl 2 switch (config) #	
Related Commands	show ib qos	
Notes		

show ib qos

show ib qos [level | match-rule | port-group | ulp]

shows the ib QoS configurations

Syntax Description	level	shows qos level configurations
	match-rule	shows qos match-rule configurations
	port-group	shows qos port-group configurations
	ulp	show qos ulp configurations
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # show ib qos level my-qos-level my-qos-level: use = my-use sl = 0 mtu-limit = 2K rate-limit = default packet-life = 0x12 pkey = -1 switch (config) #</pre>	
Related Commands		
Notes		

7.6 Subnet Manager (SM) High Availability (HA)




All nodes in an SM HA subnet must be of the same CPU type (i.e. PPC or x86).



All SM HA nodes in an SM HA cluster must run the same MLNX-OS version.

High availability (HA) refers to a system or component that is continuously operational for a desirably extended period of time.

Figure 44: SM HA Subnet



Mellanox Subnet Manager (SM) HA reduces subnet downtime and disruption as it is continuously operational for a desirably long length of time. It assures continuity of the work even when one of the SMs dies. The database is synchronized with all the nodes participating in the InfiniBand subnet and a configuration change is prepared. The synchronization is done out-of-band using an Ethernet management network.

Mellanox SM HA allows the systems' manager to enter and modify all InfiniBand SM configuration of different subnet managers from a single location. It creates an InfiniBand subnet and associates all the Mellanox management appliances that are attached to the same InfiniBand subnet into that InfiniBand subnet ID. All subnet managers can be controlled, started, or stopped from this address.

All the nodes that participate in the Mellanox SM HA are joined to the InfiniBand subnet ID and once joined, the synchronized SMs are launched. One of the nodes is elected as Master and the others are Slaves (or down). Mellanox SM HA uses an IP address (VIP) that is always directed to the SM HA master to monitor the SM state and to verify that all configurations are executed.

7.6.1 Joining, Creating or Leaving an InfiniBand Subnet ID



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 features such as Gateway HA, SM HA, and 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*].

An InfiniBand subnet is formed by a network of InfiniBand nodes interconnected via InfiniBand switches. It includes all systems that can run an SM and is part of the SM HA domain. A switch that can potentially run an SM must be a member of an InfiniBand subnet ID to be associated with the Mellanox SM HA domain. An IB subnet is recognized by its ID which is used by the system to either join or leave the subnet.

Every system that is not associated to an existing IB subnet (has never been part of an IB subnet or has left an existing one) or does not have MLNX-OS SwitchX license installed, is by default associated to a subnet called “Standalone”.

In order to create, join or leave an InfiniBand subnet, one may use the following commands:

- Create – “ib ha <IB_subnet_ID> ip <ip_addr> <netmask>”
- Join – “ib ha <IB_subnet_ID>”
- Leave – “no ib ha”



When leaving an SM HA cluster, SM configuration is not saved on the node leaving the cluster. After leaving, the configuration is reset to its default values.

For further information see [Section 7.6.5, “Creating and Adding Systems to an InfiniBand Subnet ID,”](#) on page 1506.

7.6.2 MLNX-OS Management Centralized Location

MLNX-OS centralized management infrastructure enables the user to configure or modify an existing configuration and monitor the subnet running status. MLNX-OS centralized management IP (VIP) is defined when a new subnet manager is created by running the command `ib ha <IB_subnet_ID> ip <ip_addr> <netmask>`. The created VIP is used as the current subnet master’s alias thus, assumes the same roles as the master.

The VIP always points to one of the systems part of the SM HA domain. It is always active even if one or more of the members are down. For example:

```
switch [standalone: master] (config) # ib ha subnet2 ip 192.168.10.110
255.255.255.0
switch [subnet2: master] (config) #
```

7.6.3 High Availability Node Roles

A node is an InfiniBand switch system. Every node member of an IB subnet ID has one of the following roles:

- Master – the node that manages SM configurations and provides services to the Virtual IP (VIP) addresses
- Standby – the node that replaces the Master node and takes over its responsibilities once the Master node is down
- Offline – has run an SM in the past and is currently offline, or it was created manually by the “ib smnode <node name> create” command. If the node has been removed from the environment, you can remove it from the list with the “no ib smnode xxx” command.

➤ **To see the mode of the current node, look at the CLI prompt for the following format:**

```
<host name> [<subnet ID>:<mode>] [standalone: master] (config) #
```

For example:

```
switch [ibstandalone: master] (config) #
```

To see a list of the existing nodes and details about the running state, run the command `show ib smnodes {brief}`.

7.6.4 Configuring MLNX-OS SM HA Centralized Location

The IP is used to configure or modify the existing configuration and monitor the subnet running status.

➤ **To configure the IP:**

Step 1. Enter config mode. Run:

```
switch [standalone: master] >
switch [standalone: master] > enable
switch [standalone: master] # configure terminal
```

Step 2. Configure your IP using the `ib ha <IB_subnet_ID> ip <ip_addr> <netmask>` command.

```
switch [standalone: master] (config) # ib ha subnet2 ip 192.168.10.110 255.255.255.0
switch [subnet2: master] (config) #
```

7.6.5 Creating and Adding Systems to an InfiniBand Subnet ID

➤ **To create and add systems to a subnet:**

Step 1. Log into the system from where you are creating the subnet.

Step 2. Enter config mode. Run:

```
switch [standalone: master] >
switch [standalone: master] > enable
switch [standalone: master] # configure terminal
```

Step 3. Create a new subnet using the `ib ha <IB_subnet_ID> ip <ip_addr> <netmask>` command.

```
switch [standalone: master] (config) # ib ha subnet2 ip 192.168.10.110 255.255.255.0
switch [subnet2: master] (config) #
```




You must run the `ib ha <IB_subnet_ID> ip <ip_addr> <netmask>` command **only once** per subnet ID.

Step 4. Log into the system that you are going to join to the new created subnet.

Step 5. Join the system to the subnet, using the `ib ha <IB_subnet_ID>` command.

```
switch [standalone: master] (config) # ib ha subnet2
switch [subnet2: standby] (config) #
```

7.6.6 Restoring Subnet Manager Configuration

In cases where the Subnet Manager configuration becomes corrupted or the subnet manager cannot raise any logical links it is suggested that you restore the default SM configuration.

➤ **To restore subnet manager configuration:**

Step 1. Enter config mode. Run:

```
switch [subnet2: master] > enable
switch [subnet2: master] # configure terminal
switch [subnet2: master] (config) #
```

Step 2. Run the command `ib sm reset-config`.

```
switch [subnet2: master] (config) # ib sm reset-config
```

Example:

```
switch-11a15e [SX: master] (config) # show ib smnodes brief

HA state of the switch
=====
IB Subnet HA name: SX
HA IP address:      172.30.0.12/24
Active HA nodes:    2

ID                SM-HA state    IP            SM                Priority
-----
*switch-11a15e    master          172.30.0.10    enabled - master0
switch-1133f6     standby         172.30.0.11    enabled            2

switch-11a15e [SX: master] (config) # show ib smnodes

HA state of the switch
=====
IB Subnet HA name: SX
HA IP address:      172.30.0.12/24
Active HA nodes:    2
```

```
HA node local information
Name:      switch-11a15e (active) <--- (local node)
SM-HA state: master
SM Licensed: yes
SM Running: running
SM Enabled: enabled - master
SM Priority: 0
IP:        172.30.0.10
```

```
HA node local information
Name:      switch-1133f6 (active)
SM-HA state: standby
SM Licensed: yes
SM Running: stopped
SM Enabled: enabled
SM Priority: 2
IP:        172.30.0.11
```

```
switch-11a15e [SX: master] (config) #
```



The asterisk in the example above (***switch-11a15e**) indicates the local system from where the command is running.

In order to receive information on the running state of a specific node one could run one of the following commands with its requested parameter:

- show ib smnode <name> sm-licensed
- show ib smnode <name> sm-running
- show ib smnode <name> sm-state
- show ib smnode <name> sm-priority
- show ib smnode <name> active
- show ib smnode <name> ha-state
- show ib smnode <name> ha-role

7.6.6.1 Subnet Manager Configuration

To configure the subnet manager, log into the centralized management IP (VIP). Once the SM configuration is created, the SM database is duplicated to the other nodes.



The SM **must** be configured from MLNX-OS centralized management IP (VIP). All the configurations that are not created or modified in the master node (using the VIP) are overridden by the master configuration.

The user can configure different SM parameters such as where to run the SM(s) or the SM priority by running the commands according to the desired action.

7.6.6.2 Mellanox High Availability and Opensm Handover/Failover



Mellanox Technologies products are fully compliant and interoperable with OpenSM.

Once an SM fails, the SM which takes over the subnet needs to reproduce the internal state of the failed master. Most of the information required is obtained by scanning the subnet and extracting the information from the devices. However, some information which is not stored directly in the network devices cannot be reproduced this way. InfiniBand management architecture limits such information to data exchanged between clients (either user-level programs or kernel modules) and the Subnet Administration (SA) service (attached to the SM). The SA keeps this set of client registrations in an internal data structure called SA-DB. The SA-DB information includes the multicast groups, the multicast group members, subscriptions for event forwarding and service records.

The new SM may retrieve the SA-DB by requesting the clients to re-register with the SA or by obtaining a copy of the previous master SM internal SA-DB via an SA-DB dump file. The client-re-registration offers database correctness and the SA-DB dump file replication provides lower setup time. Client re-registration is required since the SA-DB may not be up-to-date on the registrations listed in the master SM.

Furthermore, since the SM does not maintain SA-DB information for unknown nodes, it is very possible that some of the SA-DB information relating to nodes momentarily disconnected from the master SM become purged. Therefore, these nodes must re-register with the new SM when they are reconnected (they receive a client-re-register request from the SM). Relying only on client re-registration is also non-optimal as it takes some time to recreate the entire SA-DB and the network state.

Mellanox SM HA replicates the SA-DB dump file from the current master SM to all the standby SMs running on Mellanox switches. The SA-DB dump file replication provides further optimization to the standby SM that becomes master.

Standby SM loads the existing SA-DB file the old master has used. By using the existing SA-DB the amount of processing needed on client re-registration is lessened resulting in a reduced time to complete setting up the network.



SM HA does not replace InfiniBand spec requirement for client reregistration.

7.6.7 Commands

ib ha

ib ha <IB_subnet_ID> [ip <IP address> <subnet mask> [force]]
no ib ha

Creates a subnet <IB_subnet_ID> with the specified IP.
 The no form of the command removes this node from an InfiniBand subnet ID.

Syntax Description	IB subnet ID	Simple group name for shared IB config
	ip <IP address>	Assigns management IP address
	netmask	Netmask (e.g. 255.255.255.0 or /24)
	force	Joins if exists or creates if not
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # ib ha my-subnet switch (config) #	
Related Commands	show ib ha	
Notes	- A new subnet may be joined only after leaving the current one - All SM HA cluster members should be part of the same switch family (e.g. all members are part of SX60xx family or all members are part of SX65xx family)	

ib smnode

ib smnode <hostname> [create | disable | enable | sm-priority <priority>]
no ib smnode <hostname> [create | disable | enable | sm-priority]

Manages HA SM.

The no form of the command removes HA SM node configuration.

Syntax Description	hostname	Specifies <hostname> SM configuration to modify.
	create	Creates SM configuration for selected node.
	disable	Makes SM inactive on selected node.
	enable	Makes SM active on selected node.
	sm-priority <priority>	Sets SM selected node priority (0=low, 15=high).
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	switch (config) # ib smnode switch-1133ce create switch (config) #	
Related Commands	show ib smnode show ib smnodes	
Notes		

show ib smnode

show ib smnode <hostname> {active | ha-role | ha-state | ip | sm-licensed | sm-priority | sm-running | sm-state}

Displays SM High availability information.

Syntax	Description
hostname	Specifies <hostname> SM configuration to display.
active	Displays whether <hostname> is currently active.
ha-role	Displays the High Availability role of <hostname>. Possible return values are: offline, unknown, master, standby, or disabled.
ha-state	Possible return values are: offline, init, searching, joining, online, creating, waiting, leaving, join-sync, failed, removed, or regroup.
ip	Displays the local management IP address associated with the active node, <hostname>. If <hostname> is not active, the command displays “offline”.
sm-licensed	Displays if <hostname> has an SM license. The command will display “active” only if <hostname> is currently active and has a license.
sm-priority	Displays the SM priority for SM running on <hostname>.
sm-running	Displays if <hostname> has an SM running. The command will display “active” (that is, SM is running) only if <hostname> is currently active, has a license, is enabled as a potential SM, is active as SM, and if there is a maximum of 2 SMs in the fabric.
sm-state	Displays if SM is enabled to run on <hostname>.
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # show ib smnode my-hostname sm-state enabled switch (config) #</pre>

Related Commands show ib smnodes

Notes

show ib smnodes

show ib smnodes [brief]

Displays information about all the systems that are active or might be able to run SM.

Syntax Description	brief	Displays brief info on all HA nodes.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # show ib smnodes HA state of switch infiniband-default ===== IB Subnet HA name: Piranha-648-324 HA IP address: 10.7.6.238/22 Active HA nodes: 1 HA node local information Name: 43 (active) <--- (local node) SM-HA state: master SM Licensed: yes SM Running: stopped SM Enabled: disabled SM Priority: 0 IP: 10.7.7.43 HA node local information Name: 324-A (not active) SM Enabled: enabled SM Priority: 10 IP: offline switch (config) # show ib smnodes brief HA state of switch infiniband-default ===== IB Subnet HA name: Piranha-648-324 HA IP address: 10.7.6.238/22 Active HA nodes: 1 ID SM-HA state IP SM Priority ----- *43 master 10.7.7.43 disabled 0 324-A offline offline enabled 10 switch (config) #</pre>	

Related Commands `show ib smnode`

Notes

show ib ha

show ib ha [brief]

Displays information about all the systems that are active or might be able to run SM.

Syntax Description	brief	Displays HA information briefly.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # show ib ha Global HA state ===== IB Subnet HA name:subnet4 HA IP address: 192.168.10.43/24 Active HA nodes: 2 ID State Role IP SM Priority ----- switch standalone 192.168.10.42 disabled switch master 192.168.10.18 disabled switch (config) #</pre>	
Related Commands		
Notes		

7.7 Fabric Inspector

7.7.1 Running Diagnostics

➤ *To run ib fabric diagnostics:*

Step 1. Run test ib fabric to analyze the fabric.

```
switch (config) # test ib fabric
% -W- Topology file is not specified.
      Reports regarding cluster links will use direct routes.
-I- Using port 0 as the local port.
-I- Discovering ... 25 nodes (24 Switches & 1 CA-s) discovered.

-I-----
-I- Bad Guids/LIDs Info
-I-----
-I- skip option set. no report will be issued

-I-----
-I- Links With Logical State = INIT
-I-----
-W- link with LOG=INI found at direct path "24,19,17,20"
    From: a Switch   PortGUID=0x0002c90200405b98 Port=20
    To:   a Switch   PortGUID=0x0002c90200405f98 Port=18
-W- link with LOG=INI found at direct path "24,19,17,21"
    From: a Switch   PortGUID=0x0002c90200405b98 Port=21
    To:   a Switch   PortGUID=0x0002c90200405fa0 Port=18
-W- link with LOG=INI found at direct path "24,19,17,22"
    From: a Switch   PortGUID=0x0002c90200405b98 Port=22
    To:   a Switch   PortGUID=0x0002c90200405fa0 Port=17
-W- link with LOG=INI found at direct path "24,19,17,23"
    From: a Switch   PortGUID=0x0002c90200405b98 Port=23
    To:   a Switch   PortGUID=0x0002c90200405f70 Port=17
-W- link with LOG=INI found at direct path "24,19,17,24"
    From: a Switch   PortGUID=0x0002c90200405b98 Port=24
    To:   a Switch   PortGUID=0x0002c90200405f70 Port=18
-W- link with LOG=INI found at direct path "24,19,17,25"
    From: a Switch   PortGUID=0x0002c90200405b98 Port=25
    To:   a Switch   PortGUID=0x0002c90200405f80 Port=17
-W- link with LOG=INI found at direct path "24,19,17,26"
    From: a Switch   PortGUID=0x0002c90200405b98 Port=26
    To:   a Switch   PortGUID=0x0002c90200405f80 Port=18
-W- link with LOG=INI found at direct path "24,19,17,27"
    From: a Switch   PortGUID=0x0002c90200405b98 Port=27
```

```

To:  a Switch      PortGUID=0x0002c90200405f60 Port=17
-W- link with LOG=INI found at direct path "24,19,17,28"
    From: a Switch  PortGUID=0x0002c90200405b98 Port=28
    To:  a Switch  PortGUID=0x0002c90200405f60 Port=18
-W- link with LOG=INI found at direct path "24,19,17,29"
    From: a Switch  PortGUID=0x0002c90200405b98 Port=29
    To:  a Switch  PortGUID=0x0002c90200405f68 Port=17
-W- link with LOG=INI found at direct path "24,19,17,30"
    From: a Switch  PortGUID=0x0002c90200405b98 Port=30
    To:  a Switch  PortGUID=0x0002c90200405f68 Port=18
-W- link with LOG=INI found at direct path "24,19,17,31"
    From: a Switch  PortGUID=0x0002c90200405b98 Port=31
    To:  a Switch  PortGUID=0x0002c90200405f88 Port=17
-W- link with LOG=INI found at direct path "24,19,17,32"
    From: a Switch  PortGUID=0x0002c90200405b98 Port=32
    To:  a Switch  PortGUID=0x0002c90200405f88 Port=18
-W- link with LOG=INI found at direct path "24,19,17,33"
    From: a Switch  PortGUID=0x0002c90200405b98 Port=33
    To:  a Switch  PortGUID=0x0012c90200405fa9 Port=18
-W- link with LOG=INI found at direct path "24,19,17,34"
    From: a Switch  PortGUID=0x0002c90200405b98 Port=34
    To:  a Switch  PortGUID=0x0012c90200405fa9 Port=17
-W- link with LOG=INI found at direct path "24,19,17,35"
    From: a Switch  PortGUID=0x0002c90200405b98 Port=35
    To:  a Switch  PortGUID=0x0002c90200405f90 Port=17
-W- link with LOG=INI found at direct path "24,19,17,36"
    From: a Switch  PortGUID=0x0002c90200405b98 Port=36
    To:  a Switch  PortGUID=0x0002c90200405f90 Port=18

-I-----
-I- PM Counters Info
-I-----
-W- lid=0x0016 guid=0x0002c90200405a90 dev=48438 Port=23
    Performance Monitor counter      : Value
    symbol_error_counter              : 0xffff (overflow)

-I-----
-I- Fabric Partitions Report (see ibdiagnet.pkey for a full hosts list)
-I-----
-I-   PKey:0x7fff Hosts:1 full:1 partial:0

-I-----
-I- IPoIB Subnets Check
-I-----

```

```
-I- Subnet: IPv4 PKey:0x7fff QKey:0x00000b1b MTU:2048Byte rate:10Gbps SL:0x00
-W- Suboptimal rate for group. Lowest member rate:20Gbps > group-rate:10Gbps

-I-----
-I- Bad Links Info
-I- No bad link were found
-I-----

-I- Done. Run time was 511 seconds.
switch [subnet2: master] (config) #
```

Step 2. Run show fabric sm to list the subnet managers.

```
switch [subnet2: master] (config) # show fabric sm

SM - master
  Port=9 lid=0x0001 guid=0x0002c90200405f60 dev=48438 priority:0

SM - standby
  The Local Device : Port=0 lid=0x0017 guid=0x0002c9020040c6d0 dev=48438 priority:0
  Port=10 lid=0x0018 guid=0x0002c9020040b2e8 dev=48438 priority:0
switch [subnet2: master] (config) #
```

Step 3. Run show fabric pm to display the performance counters' status.

```
switch [subnet2: master] (config) # show fabric pm
% -----
Port=27 lid=0x0014 guid=0x0012c90200405a81 dev=48438
-----

symbol_error_counter = 0x0
link_error_recovery_counter = 0x0
link_down_counter = 0x0
port_rcv_errors = 0x0
port_xmit_discard = 0x0
vl15_dropped = 0x0
port_rcv_constraint_errors = 0x0
local_link_integrity_errors = 0x0
port_xmit_constraint_errors = 0x0
excessive_buffer_errors = 0x0
port_xmit_data = 0x7a1d8
port_rcv_data = 0x7a1d8
port_xmit_pkts = 0x1b23
port_rcv_pkts = 0x1b23
port_rcv_remote_physical_errors = 0x0
port_rcv_switch_relay_errors = 0x0
-----
Port=28 lid=0x0014 guid=0x0012c90200405a81 dev=48438
```

```
-----
symbol_error_counter = 0x0
link_error_recovery_counter = 0x0
link_down_counter = 0x0
port_rcv_errors = 0x0
port_xmit_discard = 0x0
vl15_dropped = 0x0
port_rcv_constraint_errors = 0x0
local_link_integrity_errors = 0x0
port_xmit_constraint_errors = 0x0
excessive_buffer_errors = 0x0
port_xmit_data = 0x7d7cf0
port_rcv_data = 0x7d7cf0
port_xmit_pkts = 0x1be2e
port_rcv_pkts = 0x1be2e
port_rcv_remote_physical_errors = 0x0
port_rcv_switch_relay_errors = 0x0
-----

Port=10 lid=0x0006 guid=0x0002c90200405f98 dev=48438
-----

symbol_error_counter = 0x0
link_error_recovery_counter = 0x0
link_down_counter = 0x0
port_rcv_errors = 0x0
...
...
...
-----

Port=26 lid=0x0014 guid=0x0012c90200405a81 dev=48438
-----

symbol_error_counter = 0x0
link_error_recovery_counter = 0x0
link_down_counter = 0x0
port_rcv_errors = 0x0
port_xmit_discard = 0x0
vl15_dropped = 0x0
port_rcv_constraint_errors = 0x0
local_link_integrity_errors = 0x0
port_xmit_constraint_errors = 0x0
excessive_buffer_errors = 0x0
port_xmit_data = 0x536d0
port_rcv_data = 0x536d0
port_xmit_pkts = 0x128a
port_rcv_pkts = 0x128a
```

```
port_rcv_remote_physical_errors = 0x0
port_rcv_switch_relay_errors = 0x0
switch [subnet2: master] (config) #
```

- Step 4.** Run `show interfaces ib` to display the status and configuration of the system's InfiniBand ports.

```
switch [subnet2: master] (config) # show interfaces ib
Slot 1 port 1 state
  Logical port state : Active
  Physical port state : 10
  Current line rate : 40.0 Gbps
  Supported speeds : 10
  Speed : 10.0 Gbps
  Supported widths : 10
  Width : 12X
  Max supported MTUs : 10
  MTU : 10
  VL capabilities : 10
  Operational VLS : 10

  RX bytes : 255
  RX packets : 255
  RX errors : 255
  Symbol errors : 255
  VL15 dropped packets: 255

  TX bytes : 255
  TX packets : 255
  TX wait : 255
  TX discarded packets: 255

Slot 1 port 2 state
  Logical port state : Active
  Physical port state : 10
  Current line rate : 40.0 Gbps
  Supported speeds : 10
  Speed : 10.0 Gbps
  Supported widths : 10
  Width : 12X
  Max supported MTUs : 10
  MTU : 10
  VL capabilities : 10
  Operational VLS : 10

  RX bytes : 255
```

```

RX packets      : 255
RX errors       : 255
Symbol errors   : 255
VL15 dropped packets: 255

```

```

TX bytes        : 255
TX packets      : 255
TX wait         : 255
TX discarded packets: 255

```

Slot 1 port 3 state

...

...

Slot 1 port 36 state

```

Logical port state : Active
Physical port state : 10
Current line rate  : 40.0 Gbps
Supported speeds   : 10
Speed              : 10.0 Gbps
Supported widths   : 10
Width              : 12X
Max supported MTUs : 10
MTU                : 10
VL capabilities    : 10
Operational VLS    : 10

```

```

RX bytes      : 255
RX packets    : 255
RX errors     : 255
Symbol errors : 255
VL15 dropped packets: 255

```

```

TX bytes      : 255
TX packets    : 255
TX wait       : 255
TX discarded packets: 255

```

switch [subnet2: master] (config) #

7.7.2 Mapping GUIDs to Node Names

To replace module GUIDs and assign meaningful names to modules use the `ib nodename` command. For further information, see *MLNX-OS CLI User Guide*.

7.7.3 Importing ibdiagnet Fabric Data

The `ib fabric import <fabric data filename>` command imports a “snapshot” of fabric data. The imported file is an output of the `ibdiagnet` tool that has previously run on any node connected to the fabric. Prior to importing fabric data, it is required to run the `ibdiagnet` tool to produce fabric data files.

Note that there are two versions of the `ibdiagnet` tool. The earlier version produces three output files describing the fabric (`ibdiagnet.db`, `ibdiagnet.pm`, and `ibdiagnet.sm`), whereas the new version produces a single file (`ibdiagnet.db_csv`).

➤ **To make an `ibdiagnet` run and import its fabric data:**

Step 1. Collect `ibdiagnet` data.

Run the following command from any node connected to the fabric. By default, `ibdiagnet` places the output file(s) under `/tmp`.

Old `ibdiagnet` version:

```
> ibdiagnet -csv -skip dup_guids zero_guids logical_state part ipoib -pm
```

New `ibdiagnet` version:

```
> ibdiagnet -pm
```

Step 2. Change directory to the `ibdiagnet` output directory.

The default directory is `/tmp`.

```
> cd <ibdiagnet output directory>
```

Step 3. Create an `ibdiagnet` output tarball. Run:

```
> tar cvzf <filename>.tgz ibdiagnet*
```

Step 4. Copy the tarball `<filename>.tgz` to the switch using the `image fetch` command. Run (in enable or config mode):

```
switch # image fetch scp://<user name>:<password>@<hostname>/<full path to <file-
name>.tgz>
100.0%[#####]
#####]
switch #
```

Step 5. Import the `ibdiagnet` file(s). Run:

```
switch [subnet2: master] (config) # ib fabric <filename>.tgz
Fabric data import successful
switch [subnet2: master] (config) #
```

7.7.4 Commands

ib fabric import

ib fabric import <filename>

Imports a “snapshot” of fabric data. It retrieves fabric data from the following ibdiagnet output files: ibdiagnet.db, ibdiagnet.sm and ibdiagnet.pm.

Syntax Description	filename	The imported file. It is an output of the ibdiagnet tool that has previously run on any node connected to the fabric, and is assumed to be a zip file with a .gz or .tgz extension.
Default	N/A	
Configuration Mode	Config	
History	3.1.1400	
Role	admin	
Example	switch (config) # ib fabric import snapshot.tgz switch (config) #	
Related Commands	show ib fabric nodes	
Notes	- To display the results of this import, you may run “show ib fabric” commands (e.g., “show ib fabric nodes type switch”) - Imported data can be displayed as long as you do not run the command “ib fabric refresh”, which overwrites the imported data - The import command cannot execute without the ibdiagnet.db file	

ib fabric monitor

ib fabric monitor
no ib fabric monitor

Enables fabric monitoring.
The no form of the command disables fabric monitoring.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.1400
Role	admin
Example	switch (config) # ib fabric monitor switch (config) # show ib fabric monitor enable switch (config) #
Related Commands	show ib fabric monitor
Notes	

ib fabric nodenames

ib fabric nodenames
no ib fabric nodenames

Imports fabric SysNames.
The no form of the command removes imported SysNames.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.1400
Role	admin
Example	switch (config) # ib fabric nodenames switch (config) #
Related Commands	
Notes	

ib fabric refresh

ib fabric refresh

Takes a “snapshot” of the current fabric data.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.1400
Role	admin
Example	switch (config) # ib fabric refresh switch (config) #
Related Commands	show ib fabric nodes
Notes	If the fabric is large, this command may take a long time to complete. this command requires license (LIC-fabric-inspector)

ib fabric transceiver-info

ib fabric transceiver-info enable
no ib fabric transceiver-info enable

Enables collection of active cable info.
 The no form of the command disables collection of active cable info.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.1400
Role	admin
Example	switch (config) # ib fabric transceiver-info enable switch (config) # show ib fabric transceiver-info enable enable switch (config) #
Related Commands	show ib fabric nodes
Notes	

test ib fabric

test ib fabric [route]

Perform infiniband fabric test

Syntax Description	route
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	monitor/admin

Example

```
switch (config) # (config) # test ib fabric
% -----
-I- Plugins load will be skipped

-----
Discovery
-I- Discovering ... 1 nodes (1 Switches & 0 CA-s) discovered.
-I- Discovery finished successfully

-I- Duplicated GUIDs detection finished successfully

-I- Duplicated Nodes Descriptions detection finished successfully

-----
Lids Check
-E- Lids Check finished with errors
-E- IBM-QA-Bay3: SX90Y3245/U1/P0 - Configured with ZERO lid

-----
Links Check
-I- Links Check finished successfully

-----
Subnet Manager
-I- SM Info retrieving finished successfully

-E- Subnet Manager Check finished with errors
-E- Not found master subnet manager in fabric

-----
Port Counters
-I- Lids Check failed, no response for some MADs can occurred
-I- Ports counters retrieving finished successfully

-I- Ports counters value Check finished successfully

-I- Ports counters Difference Check will be skipped - pause time is zero
-----
Nodes Information
-I- Lids Check failed, no response for some MADs can occurred
-W- Nodes Info retrieving finished with errors
-W- IBM-QA-Bay3: SX90Y3245/U1 - No response for MAD VSGeneralInfo

-I- FW Check finished successfully

-----
Speed / Width checks
-I- Link Speed Check (Compare to supported link speed)
-I- Links Speed Check finished successfully

-I- Link Width Check (Compare to supported link width)
-I- Links Width Check finished successfully

-----
Summary
-I- Stage           Warnings  Errors  Comment
-I- Discovery       0         0
-I- Lids Check      0         1
-I- Links Check     0         0
-I- Subnet Manager  0         1
-I- Port Counters   0         0
-I- Nodes Information 1         0
-I- Speed / Width checks 0         0
...
switch (config) #
```

Related Commands

Notes

show ib fabric connections

show ib fabric connections [attrib <speed/width>] [details] [type]

Displays the ib fabric connections with optional relevant filter.

Syntax Description	attrib <speed/width>	Attribute of connection to filter on.
	details	Displays details info.
	type	Filter connections by type. - sw-2-sw-any - Any sort of switch to switch connection - sw-2-sw-int - Internal switch to switch connection - sw-2-sw-ext - External switch to switch connection - sw-2-ca - Switch to host connection - ca-2-ca - Host to host connection
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.1.1400	
Role	admin	
Example	<pre>switch (config) # show ib fabric connections PORT-1 PORT-2 DESCRIPTION 00:08:F1:00:01:08:B5:C0-0001 00:08:F1:05:00:20:2F:7B-0035 Active 4X @ 5.0 Gbps mtu=4096 VL0 00:02:C9:03:00:61:FA:20-0001 00:08:F1:05:00:20:2F:7B-0011 Active 4X @ 10 Gbps mtu=4096 VL0, VL1 00:02:C9:03:00:61:FA:30-0002 00:08:F1:05:00:20:2F:7B-0013 Active 4X @ 10 Gbps mtu=4096 VL0, VL1 00:02:C9:03:00:61:FA:30-0001 00:08:F1:05:00:20:2F:7B-0014 Active 4X @ 10 Gbps mtu=4096 VL0, VL1 00:02:C9:03:00:5D:30:72-0004 00:08:F1:05:00:20:2F:7B-0017 Active 4X @ 10 Gbps mtu=4096 VL0 - VL7 00:02:C9:03:00:5D:30:72-0001 00:08:F1:05:00:20:2F:7B-0034 Active 4X @ 10 Gbps mtu=4096 VL0 - VL7 00:02:C9:03:00:30:95:90-0001 00:02:C9:03:00:5D:D7:B0-0003 Active 4X @ 10 (FDR10) mtu=2048 VL0 - VL7 00:02:C9:03:00:4A:E6:FE-0001 00:02:C9:03:00:5D:D7:B0-0007 Active 4X @ 10 Gbps mtu=2048 VL0 - VL7 00:02:C9:03:00:30:95:A0-0001 00:02:C9:03:00:5D:D7:B0-0008 Active 4X @ 10 (FDR10) mtu=2048 VL0 - VL7 00:02:C9:03:00:2E:E3:F0-0001 00:02:C9:03:00:5D:D7:B0-0011 Active 4X @ 10 (FDR10) mtu=2048 VL0 - VL7 switch (config) #</pre>	

Related Commands

Notes

show ib fabric messages

show ib fabric messages

Displays the InfiniBand fabric error and warning messages.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.1.1400
Role	admin
Example	<pre>switch (config) # show ib fabric messages Warning Invalid(0x02) LinkWidthSupported port 00:02:C9:03:00:30:95:90-0001 Warning Invalid(0x02) LinkWidthSupported port 00:02:C9:03:00:30:95:A0-0001 Error Internal SXX506 map error L02-19 should be S01/U1.7, not S01- 10(L02/U1.22) port 00:02:C9:03:00:49:7D:C0-0019 port 00:02:C9:03:00:5D:30:70-0010 Error Internal SXX506 map error L02-20 should be S01/U1.8, not S01- 7(L02/U1.19) port 00:02:C9:03:00:49:7D:C0-0020 port 00:02:C9:03:00:5D:30:70-0007 switch (config) #</pre>
Related Commands	
Notes	

show ib fabric monitor

show ib fabric monitor [<type>]

Displays the InfiniBand fabric monitor admin state and statistics count.

Syntax Description	type - active-links - Displays number of active point-to-point links - active-ports - Displays number of active ports in subnet - host-ports - Displays number of CA ports in subnet - nodes - Displays number of active IB chips in subnet - snapshot-time - Date/time of this snapshot - switches - Displays number of switches in subnet - systems - Displays number of active systems in subnet - unique-GUIDs - Displays total number of unique GUIDs on fabric - warnings - Displays number of topology warnings issued
Default	N/A
Configuration Mode	Any Command Mode
History	3.1.1400
Role	admin
Example	<pre>switch (config) # show ib monitor active-links 17 switch (config) # show ib monitor enable switch (config) #</pre>
Related Commands	
Notes	

show ib fabric node

show ib fabric node <system-guid> [ports]

Displays InfiniBand fabric info on one node.

Syntax Description	system-guid	The node GUID.			
	ports	Displays the info on the ports on this node.			
Default	N/A				
Configuration Mode	Any Command Mode				
History	3.1.1400				
Role	admin				
Example	<pre>switch (config) # show ib fabric node 00:02:C9:03:00:5D:D7:B0 ports System - switch node 00:02:C9:03:00:5D:D7:B0 Node details System GUID 00:02:C9:03:00:5D:D7:B0 Type SW SX60XX standalone PCI 51000:713 Ports 36 Cable support Supported PCI Device ID 51000 PCI Vendor ID 0x0002c9 Base version 1 Class version 1 Revision 161 Partition cap 8 Descriptions MF0;l-supp-SX6036: SX60XX/U1 Type Port Desc State Rate SW 00:02:C9:03:00:5D:D7:B0-0000 Switch port 0 Link Up 10 Gbps SW 00:02:C9:03:00:5D:D7:B0-0001 Port 1 Polling Up to 40 Gbps SW 00:02:C9:03:00:5D:D7:B0-0002 Port 2 Polling Up to 40 Gbps SW 00:02:C9:03:00:5D:D7:B0-0003 Port 3 Link Up 41 Gbps SW 00:02:C9:03:00:5D:D7:B0-0004 Port 4 Polling Up to 40 Gbps SW 00:02:C9:03:00:5D:D7:B0-0005 Port 5 Polling Up to 40 Gbps SW 00:02:C9:03:00:5D:D7:B0-0006 Port 6 Polling Up to 40 Gbps switch (config) #</pre>				
Related Commands					
Notes					

show ib fabric nodes

show ib fabric nodes [**cable** <cable-options>] [**role** <role-options>] [**type** <system-type>]

Displays InfiniBand fabric info on all nodes with filtering options.

Syntax Description	cable-options	Filters the list by cable type: - errors - Node with cable errors - no-errors - Node with no cable errors - supports - Node support active cables - no-support - Node does not support active cables
	role-options	Filters the list by role: - multi-chip - Systems with more than 1 nodes - single-chip - Systems with 1 node - leaf - Leaf node - spine - Spine node <system> - Any supported system
	system-type	Filters the list by system type: - switch - Switches only - host - Hosts only - router - Routers only - unknown - Unknowns systems only
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.1.1400	
Role	admin	
Example	<pre>switch (config) # show ib fabric nodes System name/GUID Type Node GUID Description 00:02:C9:03:00:5C:F7:20 SW 00:02:C9:03:00:5C:F7:20 PCI 51000:713 00:02:C9:03:00:09:DA:BD CA 00:02:C9:03:00:09:DA:BA PCI 26428:713 00:02:C9:03:00:09:28:17 CA 00:02:C9:03:00:09:28:14 PCI 26428:713 00:02:C9:03:00:5C:6E:00 SW 00:02:C9:03:00:5C:6E:00 PCI 51000:713 switch (config) #</pre>	
Related Commands		
Notes		

show ib fabric port

show ib fabric port <port-guid>

Displays InfiniBand fabric info on one port in the fabric.

Syntax Description	port-guid	The port GUID.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.1.1400	
Role	admin	
Example	<pre>switch (config) # show ib fabric port 00:02:C9:03:00:5C:6E:00-0034 SXCA07156 00:02:C9:03:00:5C:6E:00 port 00:02:C9:03:00:5C:6E:00-0034 Type SW Port state Polling Speed 2.5 Gbps Supported speeds 2.5 / 5 / 10 Gbps Width 4X Supported widths 1X, 4X Operational VLS VL0 - VL7 VL capabilities VL0 - VL7 Port GUID NA System GUID 02:C9:03:00:5C:6E:00 MTU 4096 Max supported MTUs 4096 VL arbitration high 8 VL Arbitration low 8 VL high limit 4 VL stall count 7 Has errors false Has traffic false</pre>	
Related Commands	switch (config) #	
Notes		

show ib fabric ports

show ib fabric ports [attrib <attrib-options>] [data <data-options>] [errors <errors-options>] [sm <sm-options>] [state <state-options>] [type <port-type-options>]

Displays InfiniBand fabric info on all ports with filtering options.

Syntax	Description
attrib-options	Filters the speed and width.
data-options	Filters port by data transfer counts: • none - No data • any - Any data • lots - High rate of data • little - Low rate of data
errors-options	Filters port by error counts: • none - No errors • any - Any errors • symbol - Any symbol errors • recv - Any receive errors • sym-or-recv - Any symbol or receive errors • cable - Any cable errors
sm-options	Filters port by SM running states: • active - Has an active SM • none - Does not have an SM • master - Has master SM • standby - Has a standby SM
state-options	Filters port by port state: • linkup - Link up state • polling - Polling state • unusual - Any unusual state • normal - Link up or polling state

port-type-options

Filters port by port type:

- switch-any-port - All switch ports
- switch-port0 - Switch port 0 only
- switch-not-P0 - Switch ports except 0
- switch-int - Internal switch ports
- switch-ext - External switch ports
- port-has-lid - CA or switch port 0
- has-cable-info - Port has an active cable
- has-no-cable-info - No active cable on port
- host - Host ports
- router - Router ports
- has-valid-LID - Ports with valid LIDs
- invalid-LID - Ports with invalid LIDs
- unknown - Unknown ports

Default

Configuration Mode

Any Command Mode

History

3.1.1400

Role

admin

Example

```
switch (config) # show ib fabric ports
00:02:C9:03:00:5C:F7:20 SW 00:02:C9:03:00:5C:F7:20-0000
Switch port 0 Link Up 10 Gbps
00:02:C9:03:00:5C:F7:20 SW 00:02:C9:03:00:5C:F7:20-0001 Port
1 Link Up 40 Gbps
00:02:C9:03:00:5C:F7:20 SW 00:02:C9:03:00:5C:F7:20-0002 Port
2 Polling Up to 40 Gbps
00:02:C9:03:00:5C:F7:20 SW 00:02:C9:03:00:5C:F7:20-0003 Port
3 Link Up 40 Gbps
00:02:C9:03:00:5C:F7:20 SW 00:02:C9:03:00:5C:F7:20-0004 Port
4 Polling Up to 40 Gbps
00:02:C9:03:00:5C:F7:20 SW 00:02:C9:03:00:5C:F7:20-0005 Port
5 Polling Up to 40 Gbps
00:02:C9:03:00:5C:F7:20 SW 00:02:C9:03:00:5C:F7:20-0006 Port
6 Polling Up to 40 Gbps
00:02:C9:03:00:5C:F7:20 SW 00:02:C9:03:00:5C:F7:20-0007 Port
7 Polling Up to 40 Gbps
00:02:C9:03:00:5C:F7:20 SW 00:02:C9:03:00:5C:F7:20-0008 Port
8 Polling Up to 40 Gbps
00:02:C9:03:00:5C:F7:20 SW 00:02:C9:03:00:5C:F7:20-0009 Port
9 Polling Up to 40 Gbps
00:02:C9:03:00:5C:F7:20 SW 00:02:C9:03:00:5C:F7:20-0010 Port
10 Polling Up to 40 Gbps
switch (config) #
```

Related Commands

Notes

show ib fabric system

show ib fabric system <system-guid> [nodes | ports]

Displays InfiniBand fabric info on a specific system.

Syntax Description	system-guid	The system GUID.			
	nodes	Adds list of nodes information.			
	ports	Adds list of ports information.			
Default	N/A				
Configuration Mode	Any Command Mode				
History	3.1.1400				
Role	admin				
Example	switch (config) # show ib fabric system 00:02:C9:03:00:5C:F7:20 nodes				
	System - 00:02:C9:03:00:5C:F7:20				
	Model SXCA07156				
	Revision Rev Rev 1				
	System 36 port SW				
	Element count 1				
	Description BX900S1P00355-CB5				
	Node GUID				Role
	00:02:C9:03:00:5C:F7:20				standalone
	switch (config) #				Ports 36
					Type SW
					Descripton PCI 51000:713
Related Commands					
Notes					

8 Gateway

When the network consists of two types of link protocols (Ethernet and InfiniBand), Proxy-ARP can be used to forward IPv4 packets from the Ethernet network to the InfiniBand network and vice versa. Proxy-ARP is not an IP Router, but acts as a bridge that forwards the IPoETH packets to IPoIB in Unreliable Datagram (UD).

The Proxy-ARP forwards the traffic in a single subnet.



IP Routing, InfiniBand SM and IGMP snooping must be disabled to enable Proxy-ARP.

Figure 45: Gateway



8.1 Proxy-ARP Prerequisites

Before trying to configure a Proxy-ARP in the system make sure the following conditions are met:

- Gateway license is installed (UPGR-XXXX-GW) on the switch. Run the command `show system capabilities` to verify that.



The switch system SX6036G does not require a Gateway license.

- The system profile is `vpi-single-switch`. Run the command `show system profile` to verify that. For additional information on the system profile refer to [Section 4.5, “System Profile,”](#) on page 239.
- InfiniBand and Ethernet interfaces are mapped on the system. Run the command `show ports type` to verify that. For additional information on the port mapping refer to [Section 4.4, “Virtual Protocol Interconnect \(VPI\),”](#) on page 234.
- IP routing is disabled. To disable it run:

```
switch (config)# no ip routing
```

- IGMP snooping is disabled. To disable it run:


```
switch (config)# no ip igmp snooping
```

- InfiniBand SM is disabled. To disable it run:

```
switch (config)# no ib sm
```

8.2 Proxy-ARP Overview

Figure 46: Basic Gateway Setup



8.2.1 Proxy-ARP Modes

Proxy-ARP supports unicast and multicast modes.



Multicast route configuration and learning is not available in unicast mode.



Proxy-ARP modes can be configured only when Proxy-ARP is disabled (no `ip proxy-arp`).



When Proxy-ARP mode is modified, interface configuration is removed. The user is prompted to approve this.

8.2.1.1 Proxy-ARP Unicast

An ARP request sent from an InfiniBand host is terminated at the Proxy-ARP. Then a new ARP request is generated and sent on the VLAN interface to reach the Ethernet host. The Ethernet host responds with an ARP reply to the Proxy-ARP. The Proxy-ARP then terminates it and generates a new ARP reply to the InfiniBand host. Once the destination address has been resolved, unicast traffic is passed from the InfiniBand host to the Ethernet host. The process is similar in the opposite direction (Ethernet to InfiniBand).

Figure 47: Unicast ARP Flow




Table 63 details traffic forwarding in unicast (UC) mode.

Table 63 - IPoIB Forwarding

Case	IPoETH-to-IPoIB	IPoIB-to-IPoETH
Unicast	Forward according to HA load balancing algorithm. ¹	Forward according to HA load balancing algorithm.
Multicast	Drop	Drop
Broadcast (255.255.255.255); Subnet broadcast	Forward by software in case the gateway node is the Gateway Load Balancing Protocol (GLBP) master	Forward by software in case the gateway node is the GLBP master
Master advertisement MC group (239.0.0.77); all-node group (224.0.0.1); all-router group (224.0.0.2)	Does not forward (local termination)	Does not forward (local termination)

1. See Section 8.2.3.4, “Proxy-ARP Load Balancing,” on page 1552

8.2.1.2 Proxy-ARP Multicast



Broadcast and subnet broadcast are supported in Proxy-ARP HA and are forwarded by the GLBP master.

With IB-to-ETH multicast (MC) flow, the Proxy-ARP signs up to receive InfiniBand MC notifications. Upon InfiniBand MC group creation, the Proxy-ARP is notified and as a result, the

Proxy-ARP sends a join request and triggers IGMP query. The host then receives the query and sends an IGMP join request which is forwarded to the MC router on the Ethernet side.

With ETH-to-IB MC flow, MC traffic is forwarded to the Proxy-ARP. The Proxy-ARP forwards the packet to the InfiniBand side only if a MC group is created (a MC route configured on the Proxy-ARP or dynamically according to the flow above). Otherwise, it would be dropped.



We assume that the MC router is on the Ethernet side.




Table 64 specifies how Proxy-ARP forwards multicast traffic.

Table 64 - IPoIB Multicast Forwarding

Case	IPoETH-to-IPoIB	IPoIB-to-IPoETH
Unicast	Forward	Forward
Multicast	Forward according to HA load balancing algorithm	Forward according to HA load balancing algorithm
Broadcast (255.255.255.255); Subnet broadcast; All-router group (224.0.0.2)	Forward by software in case the gateway node is the GLBP master	Forward by software in case the gateway node is the GLBP master
All-node group (224.0.0.1)	Forward in case the gateway node is the GLBP master	Forward in case the gateway node is the GLBP master

In addition to the dynamic multicast support, the software provides the ability to statically assign MC groups to a Proxy-ARP interface (by using the command “ip multicast”) and the ability to filter MC groups (by using the command “ip multicast filter”).

In case a MC group creation notification is received which matches one of the configured IP multicast filters, the Proxy-ARP ignores this notification.



If there is a contradiction between Static MC configuration and the MC filter configuration, the Static MC configuration trumps.

8.2.2 Proxy-ARP DHCP

Standard DHCP fields holding hardware (HW) addresses within DHCP message (16-byte) are not large enough to contain an IPoIB HW address (20-byte). To overcome this problem, DHCP over IB messages convey a client identifier (Option 61) field used to identify the DHCP client. For more information please refer to RFC-4390, “Dynamic Host Configuration Protocol (DHCP) over InfiniBand”.

Proxy-ARP DHCP is supported in the 2 modes described below.

8.2.2.1 DHCP Linux

DHCP Linux application on both server and client implements RFC-4390. Therefore, both driver and Proxy-ARP are not required to alter the packet in any way.

Table 65 - Proxy-ARP DHCP Linux Mode Application

DHCP Client Type	DHCP Server Type	Support
Ethernet Windows	InfiniBand Windows	+
InfiniBand Windows	Ethernet Windows	–
Ethernet Linux	InfiniBand Linux	+
InfiniBand Linux	Ethernet Linux	+
InfiniBand Linux	Ethernet Windows	+
Ethernet Linux	InfiniBand Windows	+
Ethernet Windows	InfiniBand Linux	+
InfiniBand Windows	Ethernet Linux	+

8.2.2.2 DHCP Windows



Proxy-ARP DHCP is supported with Windows driver 4.95 and later.

DHCP Windows application on both server and client does not implement RFC-4390, and hence a DHCP message is sent with ETH HW address. Therefore, on host egress traffic, IPoIB Windows driver alters the packets and translates the ETH (6-byte) address according to RFC-4390 and vice versa.

A DHCP packet from IB to ETH is transmitted from the host in the format of RFC-4390, and is received on the ETH side as such. As Windows DHCP application does not implement RFC-4390, the packet is dropped. Therefore, Proxy-ARP is required to translate RFC-4390 into ETH (6-byte) address.

Table 66 - Proxy-ARP DHCP Windows Mode Application

DHCP Client Type	DHCP Server Type	Support
Ethernet Windows	InfiniBand Windows	+
InfiniBand Windows	Ethernet Windows	+
Ethernet Linux	InfiniBand Linux	—
InfiniBand Linux	Ethernet Linux	—
InfiniBand Linux	Ethernet Windows	+
Ethernet Linux	InfiniBand Windows	—
Ethernet Windows	InfiniBand Linux	—
InfiniBand Windows	Ethernet Linux	—

8.2.2.3 Configuring Proxy-ARP DHCP

➤ **To configure DHCP on Proxy-ARP:**

Step 1. Enable Proxy-ARP. Run:

```
switch (config)# ip proxy-arp
```

Step 2. (Optional) Configure Proxy-ARP DHCP mode. Run:

```
switch (config) # proxy-arp dhcp linux
```

Step 3. Verify the configured Proxy-ARP DHCP mode. Run:

```
switch (config) # show proxy-arp dhcp
Proxy-arp dhcp: linux
switch (config) #
```

8.2.3 Proxy-ARP High Availability



All nodes in a Proxy-ARP HA group must feature the same system profile, CPU type (i.e. PPC or x86), and Proxy-ARP mode. Any mismatch may cause nodes not to join a cluster and will lead to errors in the log.



Is it recommended to have all nodes in a Proxy-ARP HA group configured with the same system resource table.




Proxy-ARP interfaces support IGMPv2 only.



Only the GLBP master handles Proxy-ARP DHCP when it is Proxy-ARP high availability is active.

High availability (HA) refers to a system or component that is continuously operational for a desirably extended period of time. The following sections introduce Mellanox high availability Proxy-ARP.

Figure 49: High Availability Proxy-ARP Interface



All gateway nodes must be on the same management IP subnet.

Proxy-ARP HA reduces downtime and disruption as it is continuously operational for a desirably long length of time. It assures continuity of the work even when one Gateway fails. The database is synchronized with all the nodes participating in the Proxy-ARP group and a configuration change is prepared. The synchronization is done out-of-band using an Ethernet management network.

All the nodes that participate in the Proxy-ARP group are joined under a Proxy-ARP group name. One of the nodes is elected as master and the others become slaves. Proxy-ARP HA uses an IP address (VIP) that is always directed to the master node. The configuration is set via the VIP to the master node which in turn distributes it to the slave nodes.



The VIP address must be on the same management IP subnet.

8.2.3.1 Joining, Creating or Leaving a Proxy-ARP HA Group



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 features such as Gateway HA, SM HA, and 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*].

A Proxy-ARP group is formed by a network of Gateway nodes. Each group may be composed of up to 16 nodes. A Proxy-ARP group has a name identifier used for joining or leaving the group.

To create, join or leave a Gateway subnet, one may use the following commands:

- Create – proxy-arp ha <group-name> ip <ip_addr> <mask>
- Join – proxy-arp ha <group-name>
- Leave – no proxy-arp ha

For further information see [Section 8.2.3.6, “Managing High Availability Proxy-ARP Groups,” on page 1553.](#)

8.2.3.2 MLNX-OS Management Centralized Location

MLNX-OS centralized management infrastructure enables the user to configure or modify an existing configuration and monitor a Proxy-ARP group’s running status. MLNX-OS centralized management IP (VIP) is defined when the first Proxy-ARP HA group member is added by using the command proxy-arp ha <group-name> ip <ip_addr> <mask>. Proxy-ARP HA always directs the configured IP address (VIP) to the master node.

The VIP is always serviced by the node which is currently the master of the HA group. It is always active even if one or more of the members are down.

8.2.3.3 High Availability Node Roles

A node, in this case, is a Gateway switch system. Every node member of a Proxy-ARP HA group has one of the following roles:

- Master – the node that manages HA configurations and service to the VIP addresses.
- Standby – each node that can replace the master node and take over its responsibilities once the master node is down

Nodes support the following functions:

- Discovery – network discovery
- Master election – electing a master node among the nodes running in a group.
- Failure detection – detecting failure if a node goes down. In case of failure the system recovers and ensures a master node is available.
- Configuration replication – configuration on the master node is replicated on the slave node.

➤ **To see the mode of the current node, look at the CLI prompt for the following format:**

```
<host name> [<subnet ID>:<mode>] (config) #
```

To see a list of the existing nodes and details about the running state, run the command `show proxy-arp ha`.

8.2.3.4 Proxy-ARP Load Balancing



Changing gateway node mode from standalone to HA for the first time may cause traffic loss for up to 1 minute.



Changing modes from HA to standalone removes member related configuration.

Proxy-ARP HA features two load balancing modes:

- Active-Active – the Gateway nodes within the group are all active and share the traffic load
- Active-Standby – one Gateway nodes handles all the traffic while the others are on standby mode

The Active-Active load balancing algorithm is based on a VRRP-like protocol. IP addresses of hosts are distributed between the active Gateway systems. Upon adding or removing a new member the load is redistributed among the members accordingly. To minimize traffic loss, when a member is added, traffic is passed to it only once the new member is operationally ready. Once the new member begins serving a group of IP hosts, it sends gratuitous ARPs for the purpose of rapid stabilization. Other than in cases when a Proxy-ARP HA member leaves or joins the group, traffic does not move between existing members so as to minimize loss of traffic.

Active-Active load balancing is of three types:

- InfiniBand base IP – based on the destination IP for Ethernet sourced traffic, and source IP for InfiniBand sourced traffic
- Ethernet base IP – based on the destination IP for InfiniBand sourced traffic, and source IP for Ethernet sourced traffic
- Source IP Destination IP – best distribution of traffic when many hosts are populated on both sides

In Active-Standby load balancing the Proxy-ARP HA member with the highest priority is the one selected to pass the traffic. However, if two members have the same priority, the one with the highest GUID number carries all the traffic.

Member nodes in the Proxy-ARP HA group are organized into a virtual ring. Within this group, the following fail scenarios may occur depending on the failing node:

- Master fails – the standby detects it, becomes master, and selects a new standby node from the list of members making failover fast

- Standby fails – the master selects a new standby out of the remaining members

8.2.3.5 Prerequisites

The HA Proxy-ARP prerequisites match those of the standalone Proxy-ARP prerequisites. Refer to [Section 8.1, “Proxy-ARP Prerequisites,”](#) on page 1544.

8.2.3.6 Managing High Availability Proxy-ARP Groups

➤ *To create and add systems to a subnet:*

Step 1. Log into the system from where you are creating the subnet.

Step 2. Create a new Proxy-ARP group using the command `proxy-arp ha <group-name> ip <ip_addr> <mask>`. Run:

```
GatewayA [standalone: master] (config) # proxy-arp ha my-group ip 10.10.10.10 /24
GatewayA [my-group: master] (config) #
```



You must run the command `proxy-arp ha <group-name> ip <ip_addr> <mask>` **only once** per Proxy-ARP group.



The group IP is used to as a single point for gateway configuration and monitoring the subnet’s running status (by running “show” commands).

Step 3. Log into the system that you are going to join to the new created Proxy-ARP group.

Step 4. Join another system to the Proxy-ARP group, using the command `proxy-arp ha <group-name>`. Run:

```
GatewayB [standalone: master] (config) # proxy-arp ha my-group
GatewayB[my-group: standby] (config) #
```

Once a Proxy-ARP HA group is created, the first member maintains its configuration. Configuration of the Proxy-ARP HA may then be performed to suit the subnet needs. When new nodes join the group, the configuration is duplicated to them.


A new node joining a group should serve all the interfaces. For each interface addition, set the Member IP addresses using the command `interface proxy-arp <pra-id> ha member <hostname> ip address <ip-address>`.

Removing a member from an active Proxy-ARP group must be done using the command `no proxy-arp ha` from the gateway management interface.

8.2.4 Proxy-ARP Interface

Proxy-ARP is a logical interface which owns the IP address, VLAN and PKEY values of the subnet.

Figure 50: Proxy-ARP Interface



The Proxy-ARP interface has the following attributes:

Table 67 - Proxy-ARP Interface Attributes

Parameter	Description	Scope	Access
Admin state	Enables or disables the interface	Interface	RW
Operational state	Shows the operational state of the interface (up or down)	Interface member	RO
Description	Sets an alias text to the interface	Interface	RW
IP address	IP address	Interface member	RW
IP netmask	Network prefix mask	Interface	RW
GUID	Interface GUID	Interface member	RO
MTU	Maximum transmission unit	Interface	RW
VLAN	VLAN value	Interface	RW
PKEY	PKEY value	Interface	RW
Counters admin	Enables a set of counters on the interface	Interface	RW
Set of counters	Interface counters	Interface member	RO
Priority	Interface member priority	Interface member	RW
HA timers	Timers the control Proxy-ARP HA behavior	Interface member	RW

Table 67 - Proxy-ARP Interface Attributes

Parameter	Description	Scope	Access
Multicast routes	Multicast	Interface	RW
Multicast filters	Multicast	Interface	RW

8.2.5 Proxy-ARP HA Resources

8.2.5.1 HA Active-Standby Mode

The resources available for Proxy-ARP at HA Active-Standby load balancing are as single node. See [Section 8.2.5, “Proxy-ARP HA Resources,” on page 1555](#) for more information.

8.2.5.2 HA Active-Active IB Based IP Mode

The resources available for Proxy-ARP as HA IB based IP mode with Active-Active load balancing are as follows:

- Number of InfiniBand ARP entries – the number of InfiniBand ARP entries per node (please see [Section 8.2.5, “Proxy-ARP HA Resources,” on page 1555](#)) multiplied by the number of nodes in the Gateway HA Group
- Number of InfiniBand GRH ARP entries – the number of InfiniBand GRH ARP entries per node (please see [Section 8.2.5, “Proxy-ARP HA Resources,” on page 1555](#)) multiplied by the number of nodes in the Gateway HA Group
- Number of MC routes – the max MC routes per node (please see [Section 8.2.5, “Proxy-ARP HA Resources,” on page 1555](#)) multiplied by the number of nodes in the Gateway HA group, 3072.



Due to a resource limitation in the number of sockets that can be opened, the number of static MC groups that can be configured while in a cluster has been limited to 600. In other words, if more than 600 static MC groups are configured on the switch, and the user attempts to join a cluster, the operation fails.

8.3 Proxy-ARP Event Notifications

Table 68 - Supported Proxy-ARP Event Notifications

Scenario	Notes	Event Message
Operational state of the interface goes up	Severity level: Notice	Interface proxy-arp <instance> changed state to UP
Operational state of the interface goes down	Severity level: Notice	Interface proxy-arp <instance> changed state to DOWN
Proxy-ARP HA master detects a new member has been added to the Proxy-ARP interface	Severity level: Notice	Interface proxy-arp <instance> detected added member <member>
Proxy-ARP HA master detects a member has been removed from the Proxy-ARP interface	Severity level: Notice	Interface proxy-arp <instance> detected removed member <member>
Member is elected Proxy-ARP HA master	This message appears on the member entering master state. Severity level: Notice	Interface proxy-arp <instance> has changed to master state
Member is no longer Proxy-ARP HA master	This message appears on the member leaving master state. Severity level: Notice	Interface proxy-arp <instance> has left master state
Rate at which keepalive packets are lost exceeds configured threshold	This message appears on the interface master only. Severity level: Warning	Interface proxy-arp <instance> lost KA messages from member <member>, exceeds threshold
Rate at which keepalive packets are lost falls below configured threshold	This message appears on the interface master only. Severity level: Notice	Interface proxy-arp <instance> resumed receiving KA messages from member <member>
Rate at which keepalive packets are lost on VLAN/PKEY leg exceeds configured threshold.	This message appears on the interface master only. Severity level: Warning	Interface proxy-arp <instance> lost KA messages from member <member> on <vlan/pkey> leg, exceeds threshold
Rate at which keepalive packets are lost on VLAN/PKEY falls below configured threshold	This message appears on the interface master only. Severity level: Notice	Interface proxy-arp <instance> resumed receiving KA messages from member <member> on <vlan/pkey> leg

Table 68 - Supported Proxy-ARP Event Notifications

Scenario	Notes	Event Message
Advertisement packets are lost on VLAN/PKEY leg of backup member	This message appears on the active backup member which lost the packets. Severity level: Warning	Interface proxy-arp <instance> not receiving advertisement messages on <vlan/pkey> leg
Advertisement packets are recovered on VLAN/PKEY leg of backup member	This message appears on the active backup member which lost the packets. Severity level: Notice	Interface proxy-arp <instance> resumed receiving advertisement messages on <vlan/pkey> leg

8.4 Proxy-ARP Configuration

8.4.1 Proxy-ARP Mode Configuration

➤ *To transition from unicast to multicast:*

```
switch [standalone: master] (config) # proxy-arp mode multicast

Warning! All proxy-arp interfaces configuration is going to be deleted.
Type 'yes' to confirm mode change: yes

switch [standalone: master] (config) # show proxy-arp mode
Proxy-arp mode: multicast
```

➤ *To transition from multicast to unicast:*

```
switch [standalone: master] (config) # proxy-arp mode unicast

Warning! All proxy-arp interfaces configuration is going to be deleted.
Type 'yes' to confirm mode change: yes

switch [standalone: master] (config) # show proxy-arp mode
Proxy-arp mode: unicast
```

8.4.2 Standalone Proxy-ARP Configuration

➤ *To configure Proxy-ARP in the system:*

Step 1. Make sure the prerequisites conditions are met. Verify that gateway is supported as part of the system capabilities. Run:

```
switch [standalone: master] (config)# show system capabilities
IB: Supported
Ethernet: L3
GW: Supported
Max SM nodes:648
Ethernet Max licensed speed: 40Gbps
IB max licensed speed: FDR
switch [standalone: master] (config)#
```

Step 2. Set Proxy-ARP mode. Run:

```
switch [standalone: master] (config)# proxy-arp mode multicast
```



Proxy-ARP mode is unicast by default. This step should only be performed if you wish to configure multicast routes.

Step 3. Enable Proxy-ARP. Run:

```
switch [standalone: master] (config)# ip proxy-arp
```

Step 4. Set system resource mode. Run:

```
switch [standalone: master] (config)# system resource table loose
```



System resource table is strict by default. This step should only be performed if you wish to move to loose resource table mode.



The user is prompted to reload the system. If approved, system configuration is saved prior to reload.

Step 5. Create a Proxy-ARP interface. Run:

```
switch [standalone: master] (config)# interface proxy-arp 1
switch [standalone: master] (config interface proxy-arp 1)#
```

Step 6. Set an IP address and network mask to the Proxy-ARP interface. Run:

```
switch [standalone: master] (config interface proxy-arp 1)# ip address
20.20.20.20
switch [standalone: master] (config interface proxy-arp 1)# ip netmask /24
```



Each Proxy-ARP interface must have a unique IP subnet.



The Proxy-ARP interface IP must be on a different IP subnet than that of the management interfaces.

Step 7. Create a VLAN. Run:

```
switch [standalone: master] (config)# vlan 10
switch [standalone: master] (config vlan 10)# exit
```

Step 8. Add a VLAN to the interface. Run:

```
switch [standalone: master](config interface proxy-arp 1)# ip vlan 10
```

Step 9. Add a PKEY to the interface. Run:

```
switch [standalone: master](config interface proxy-arp 1)# ip pkey 0x7fff
```



Each Proxy-ARP interface must have a unique VLAN/PKEY pair.

Step 10. Enable the Proxy-ARP interface. Run:

```
switch [standalone: master] (config interface proxy-arp 1)# no shutdown
```

- Step 11.** Make sure one of the Ethernet or port-channel ports are configured with VLAN 10 and that the Proxy-ARP part of the VLAN is operationally up. For example:

```
switch [standalone: master] (config interface ethernet 1/1)# switchport
access vlan 10
switch [standalone: master] (config interface ethernet 1/1)# show interfaces
proxy-arp 1
Proxy-arp 1
  Admin state: Enabled
  Operational state: Up
  GUID: 00:02:C9:03:00:66:08:63
  Internet Address: 20.20.20.20/24
  Broadcast Address: 20.20.20.255
  Description: N/A
  MTU: 1500
  Counters: Disabled
  Member interfaces: vlan 10, pkey 0x7fff
switch (config)#
```

- Step 12.** (Optional) Configure a route to the default gateway in the subnet. Run:

```
switch [standalone: master] (config interface proxy-arp 1)# ip route default
20.20.20.254
```



The default gateway configuration is not used for management purposes.

- Step 13.** (Optional) Configure a multicast route. While in Proxy-ARP multicast mode, run:

```
switch [standalone: master] (config interface proxy-arp 1)# ip multicast
237.0.1.5
```

➤ **To verify the Proxy-ARP configuration:**

- Step 1.** Display the Proxy-ARP interface configuration. Run:

```
switch [standalone: master] (config)# show interfaces proxy-arp 1
Proxy-arp 1
  Admin state: Enabled
  Operational state: Up
  GUID: 00:02:C9:03:00:66:08:63
  Internet Address: 20.20.20.20/24
  Broadcast Address: 20.20.20.255
  Description: N/A
  MTU: 1500
  Counters: Disabled
  Member interfaces: vlan 10, pkey 0x7fff
switch [standalone: master] (config)#
```

Step 2. Display the Proxy-ARP brief status. Run:

```
switch [standalone: master] (config)# show interfaces proxy-arp brief
Interface      Description      State Member interfaces
-----
Proxy-arp 1 N/A      Up      vlan 10, pkey 0x7fff
switch [standalone: master] (config)#
```

Step 3. Display the routing table. Run:

```
switch [standalone: master] (config) # show ip route interface proxy-arp 1
Destination Mask      Gateway      Interface      Source      Distance/Metric
20.20.20.0  255.255.255.0  0.0.0.0      proxy-arp 1  kernel      0/0
default      0.0.0.0      20.20.20.254 proxy-arp 1  static      0/0
switch [standalone: master] (config) #
```

Step 4. Display the multicast routing table. Run:

```
switch [standalone: master] (config) # show ip multicast interface proxy-arp 1
Proxy-arp 1 multicast list:

Total number of entries: 5

Address      Interface      Source
-----
224.0.0.1    proxy-arp 2    Dynamic
224.0.0.2    proxy-arp 2    Dynamic
237.0.1.5    proxy-arp 2    Static
239.0.0.77   proxy-arp 2    Dynamic
255.255.255.255 proxy-arp 2    Dynamic
switch [standalone: master] (config) #
```

8.4.3 High Availability Proxy-ARP Configuration

➤ *Configure the following on GatewayA:*

```
# ssh admin@10.10.10.11
GatewayA [standalone: master] (config)# vlan 10
GatewayA [standalone: master] (config vlan 10)# exit
GatewayA [standalone: master] (config)# interface ethernet 1/1 switchport access
vlan 10
GatewayA [standalone: master] (config)# proxy-arp mode unicast force
GatewayA [standalone: master] (config)# ip proxy-arp
GatewayA [standalone: master] (config)# system resource table loose \\*optional*\\
GatewayA [standalone: master] (config)# proxy-arp ha my-group ip 10.10.10.10
255.255.255.0
GatewayA [my-group: master] (config)
```

➤ **Configure the following on GatewayB:**

```
# ssh admin@10.10.10.12
GatewayB [standalone: master](config)# vlan 10
GatewayB [standalone: master] switch (config vlan 10)# exit
GatewayB [standalone: master](config)# interface ethernet 1/1 switchport access
vlan 10
GatewayB [standalone: master](config)# proxy-arp mode unicast force
GatewayB [standalone: master](config)# ip proxy-arp
GatewayA [standalone: master] (config)# system resource table loose \\*optional*\\
GatewayB [standalone: master](config)# proxy-arp ha my-group
GatewayB [my-group: standby] (config)
```



The user must make sure all gateway nodes are running with the same Proxy-ARP mode (multicast or unicast).



The user must make sure all gateway nodes are running with the same system resource table mode (strict or loose).

➤ **Configure the following on the VIP:**

```
# ssh admin@10.10.10.10
GatewayA [my-group: master] (config)# interface proxy-arp 1
GatewayA [my-group: master] (config interface proxy-arp 1)# ip netmask /24
GatewayA [my-group: master] (config interface proxy-arp 1)# ip vlan 10
GatewayA [my-group: master] (config interface proxy-arp 1)# ip pkey 0x7fff
GatewayA [my-group: master] (config interface proxy-arp 1)# ha member Gate-
wayA ip address 20.20.20.20
GatewayA [my-group: master] (config interface proxy-arp 1)# ha member Gate-
wayB ip address 20.20.20.21
GatewayA [my-group: master] (config interface proxy-arp 1)# no shutdown
GatewayA [my-group: master] (config interface proxy-arp 1)# ip route default
10.10.10.250
```



The IP address of the HA member must not be the IP of the management interface. It should be on a different IP subnet. The HA group IP, however, should be the same IP subnet of the management IP of the system.

➤ **To verify the Proxy-ARP configuration on the VIP:**

```
GatewayA [my-group: master] (config) # show interfaces proxy-arp 1 ha

Proxy-arp 1
  Keep Alive loss threshold: 1
  Keep Alive loss interval: 120 seconds
  Load balancing algorithm: ib-base-ip
  IP masklen: 24
  Admin state: Enabled
  MTU: 1500
  Counters: Disabled
  Bridged interfaces: vlan 10, pkey 0x7fff
  Number of members: 2
  Hostname  Description  Admin State  LB State  Operational State  IP  Priority
  -----
  --
  GatewayA  host1-dscr  Enabled      Active    Up                 20.20.20.20  100
  GatewayB  host2-dscr  Enabled      Active    Up                 20.20.20.21  100
GatewayA [my-group: master] (config) # show proxy-arp ha
Load balancing: ib-base-ip
Number of Proxy-ARP interfaces: 1

Proxy Arp VIP:
=====
Proxy-arp group name: my-group
HA VIP address: 10.10.10.10/24
Active nodes: 2
Hostname  State  IP Address
-----
GatewayA  master  10.10.10.11
GatewayB  standby 10.10.10.12
GatewayA [my-group: master] (config) #
```

➤ **To change the hostname in an active Proxy-ARP HA cluster:**

- Step 1.** Remove Proxy-ARP member configuration related to the switch with the hostname to be renamed.
- Step 2.** Add Proxy-ARP member configuration related to the switch with the new hostname to be configured.
- Step 3.** Wait until load balancing is distributed to all members but this switch.
- Step 4.** Change the switch hostname.
- Step 5.** Save configuration.
- Step 6.** Verify load balancing is distributed to all members.

For example, assuming GatewayA and GatewayB are members of “my-group” and we wish to change GatewayB’s name to GatewayC.

➤ **To create the setup:**

Step 1. Connect via interface management of GatewayA. Run:

```
GatewayA [standalone: master] (config) # proxy-arp mode unicast force
GatewayA [standalone: master] (config) # ip proxy-arp
GatewayA [standalone: master] (config) # proxy-arp ha my-group ip 10.10.10.10 /24
GatewayA [my-group: master] (config) #
```

Step 2. Connect via interface management of GatewayB. Run:

```
GatewayB [standalone: master] (config) # proxy-arp mode unicast force
GatewayB [standalone: master] (config) # ip proxy-arp
GatewayB [standalone: master] (config) # proxy-arp ha my-group
GatewayB [my-group: standby] (config) #
```

Step 3. Connect via VIP. Run:

```
GatewayA [my-group: master] (config)# interface proxy-arp 1 netmask /24
GatewayA [my-group: master] (config)# interface proxy-arp 1 ha member GatewayA ip
20.20.20.20
GatewayA [my-group: master] (config)# interface proxy-arp 1 ha member GatewayB ip
20.20.20.21
GatewayA [my-group: master] (config)# no interface proxy-arp 1 shutdown
```

➤ **To change the hostname from GatewayB to GatewayC:**

Step 1. Delete GatewayB member configuration. Connect via VIP. Run:

```
GatewayA [my-group: master] (config) # no interface proxy-arp 1 ha member GatewayB
```

Step 2. Configure GatewayC member configuration. Connect via VIP. Run:

```
GatewayA [my-group: master] (config) # interface proxy-arp 1 ha member GatewayC ip
20.20.20.22
```

Step 3. Connect through the GatewayB management interface IP, and change the hostname. Run:

```
GatewayB [my-group: standby] (config) # hostname GatewayC
GatewayC [my-group: standby] (config) #
```

8.5 Advanced Settings

8.5.1 Default Gateway

It is recommended to configure a route to the default gateway in the subnet. If the default gateway is not configured, unregistered unicast traffic is dropped.

8.5.2 vTCA Interface

A virtual Target Channel Adapter (vTCA) is an end-point of InfiniBand fabric. Gateway needs a vTCA enabled on the switch in order to function (SMA port #37).

The vTCA is available in VPI single switch mode and when Proxy ARP is enabled.

The vTCA interface is enabled by default. However, if the SM disables this interface, it can be re-enabled by running the following command:

```
switch (config)# no sma port 1 shutdown
switch (config)# show sma port 1
Enabled
switch (config)#
```

When using InfiniBand tools such as `iblinkinfo`, `smpquery`, or `ibnetdiscover` the user is able to see the status of the vTCA interface.

```
# iblinkinfo
...
6 37[]==(4X 14.0625 Gbps Active/LinkUp)==> 7 1[] "Mellanox vTCA switch-
626a54" ( )
...
#
# smpquery -D pi 0 1 37
Port info: DR path slid 65535; dlid 65535; 0 port 1
...
CapMask:.....0x251486a
          IsSM
          IsTrapSupported
          IsAutomaticMigrationSupported
          IsSLMappingSupported

          IsSystemImageGUIDsupported
          IsExtendedSpeedsSupported
          IsCommunicationManagementSupported
          IsVendorClassSupported
          IsCapabilityMaskNoticeSupported
          IsClientRegistrationSupported

...
...
LinkState:.....Active
PhysLinkState:.....LinkUp
...
...
#
```

```
# ibnetdiscover
#
# Topology file: generated on Tue Jan 29 15:08:32 2013
#
# Initiated from node 0002c903003531b0 port 0002c903003531b1

...
...
Ca      1 "H-0002c903006cc4f2"          # "Mellanox vTCA switch-626a54"
[1](2c903006cc4f2)      "S-0002c903006cc4f1"[37]          # lid 7 lmc 0
"MF0;switch-626a54: SX1036/U1" lid 6 4xFDR

vendid=0x2c9
devid=0x1003
sysimgguid=0x2c90300431cd3
caguid=0x2c90300431cd0
...
...

#
```

8.5.3 MTU

Make sure that the InfiniBand subnet MTU is similar to the Ethernet subnet MTU. In most cases the default the MTU is 1500 bytes for Ethernet subnets while 4K in InfiniBand. Run:

```
switch (config)# interface ethernet 1/1 mtu 4000
switch (config)# interface ib 1/10 mtu 4000
switch (config)# interface proxy-arp 1 mtu 4000
```



Proxy-ARP MTU is limited to 4092.

8.6 Commands

8.6.1 Config

ip proxy-arp

ip proxy-arp
no ip proxy-arp

Enables Proxy-ARP feature.
 The no form of the command disables Proxy-ARP feature.

Syntax Description	N/A
Default	IP Proxy-ARP is disabled
Configuration Mode	Config
History	3.3.0000
Role	admin
Example	switch (config) # ip proxy-arp switch (config) #
Related Commands	show ip proxy-arp
Note	- The command “ip proxy arp” can be enabled only if the following features are disabled: •ip routing •ip igmp snooping •ib sm - When IP Proxy-ARP is disabled the user is still able to configure Proxy-ARP but it only takes affect when Proxy-ARP is enabled. The no form of the command does not delete Proxy-ARP configuration. - This command is not available on VIP. - The no form of the command is not available when the gateway node is in an HA group. The user should first leave the group (“no proxy-arp ha <my-group>”) and then use the no form of the command.

proxy-arp mode

proxy-arp mode {unicast | multicast} [force]
no proxy-arp mode [force]

Configures IP Proxy-ARP to work in either unicast or multicast mode.
 The no form of the command resets this parameter to its default.

Syntax Description	unicast	Configures unicast mode
	multicast	Configures multicast mode
	force	Forces configuration on the system
Default	unicast	
Configuration Mode	Config	
History	3.4.0000	
Role	admin	
Example	switch (config) # proxy-arp mode unicast switch (config) # show proxy-arp mode Proxy-arp mode: unicast	
Related Commands		
Note	• The difference between unicast and multicast mode is in the hardware resource allocation. • The command is only available when IP Proxy-ARP is disabled. • The command is only available from the gateway node IP. • The force parameter in the no form of the command bypasses the confirmation prompt • All Proxy-ARP configurations is deleted upon Proxy-ARP mode change. The user is prompted with a warning message informing him of this, and is required to confirm the command.	

proxy-arp dhcp

proxy-arp dhcp {linux | windows}
no proxy-arp dhcp

Configures Proxy-ARP DHCP mode.
The no form of the command resets this parameter to its default.

Syntax Description	linux	Sets DHCP mode to Linux
	windows	Sets DHCP mode to Windows
Default	Linux	
Configuration Mode	Config	
History	3.4.2100	
Role	admin	
Example	switch (config) # proxy-arp dhcp linux	
Related Commands		
Note	• If Proxy-ARP HA is disabled, the command is available from the box management IP address • If Proxy-ARP HA is enabled, the command is available from the VIP address • If the command is run on the Proxy-ARP HA master, this configuration is replicated to all cluster nodes	

interface proxy-arp

interface proxy-arp <pra-id>[-<pra-id>]
no interface proxy-arp <pra-id>[-<pra-id>]

Creates a Proxy-ARP interface, and enters a new configuration mode.
 The no form of the command deletes the interface configuration.

Syntax Description	pra-id	A Proxy-ARP interface ID. Range: 1-128.
Default	N/A	
Configuration Mode	Config	
History	3.3.0000	First version
	3.3.4302	Updated Proxy-ARP interface limitation in Note field
	3.3.4402	Added HA note
	3.4.0000	Added range option
Role	admin	
Example	<pre>switch (config) # interface proxy-arp 5-7 switch (config interface proxy-arp 5-7) # exit switch (config) # interface proxy-arp 1 switch (config interface proxy-arp 1) #</pre>	
Related Commands	<pre>show interfaces proxy-arp show interfaces proxy-arp brief</pre>	
Note	- Up to 32 Proxy-ARP interfaces can be created. - This command is only available on the master (standalone or Proxy-ARP HA group master) and via VIP. - When in the configuration mode of multiple Proxy-ARP interfaces, any configuration performed is applied on all interfaces within that range.	

proxy-arp ha

proxy-arp ha <group-name> [ip <ip-address> <ipv-mask> [force]]
no proxy-arp ha

Creates a Proxy-ARP group or maps the system to an existing group.
The no form of the command removes the system from the Proxy-ARP group.

Syntax Description	group-name	64B string of the Proxy-ARP group name.
	ip <ip-address> <ipv-mask>	Sets the VIP and mask of the Proxy-ARP group. The mask is IPv4 and can be in the format of “255.255.255.0” or “/24”.
	force	Forces the Proxy-ARP IP configuration when executing from a group member that is not the master.
Default	N/A	
Configuration Mode	Config	
History	3.3.4402	First version
	3.4.0000	Updated Note section
Role	admin	
Example	switch (config) # proxy-arp ha my-group ip 10.10.10.10 /24	
Related Commands	show proxy-arp ha	
Note	- The VIP must be in the management interface’s subnet. - The command fails if the switch is already a member of a Proxy-ARP group, or if there is not at least one active member of the Proxy-ARP group currently active on the management network - Once configuring this feature, the user should connect to the switch via this VIP to configure Proxy-ARP - This command is applicable via management gateway node IP only - This command is only available when Proxy-ARP mode is unicast - Leaving a Proxy-ARP group prompts the message: “Member <member-name> has been removed from all configured proxy-arp interfaces”	

proxy-arp ha lb-algorithm

proxy-arp ha lb-algorithm <type>
no proxy-arp ha lb-algorithm

Configures the load balancing algorithm.
 The no form of the command resets the value to its default value.

Syntax Description	type	Possible load-balancing algorithms. Possible values are: - activity-standby – active/standby HA scheme. Only one switch passes traffic. - ib-base-ip – load balancing based on the InfiniBand host's IP address (source or destination)
Default	ib-base-ip	
Configuration Mode	Config	
History	3.3.4402	
Role	admin	
Example	<pre>switch (config) # proxy-arp ha lb-algorithm ib-base-ip switch (config) #</pre>	
Related Commands	show proxy-arp ha	
Note	This command is only available on the Proxy-ARP HA group master and via VIP.	

sma port

sma port <number> shutdown
no sma port <number> shutdown

Disables the SMA on virtual TCA.
The no form of the command enables the SMA on virtual TCA.

Syntax Description	number	SMA port number.
Default	SMA port 1 is enabled.	
Configuration Mode	Config	
History	3.3.0000	
Role	admin	
Example	switch (config) # sma port 1 shutdown switch (config) #	
Related Commands	show sma port	
Note	- This port must be enabled for gateway functionality. In case the SM disabled this port via MAD interface, this command provides the administrator the ability to re-enable it back (without the need to reload the switch). - SMA ports are currently limited to only one port - In regular operation there is no need to change the configuration of this parameter - This command is applicable via management gateway node IP only.	

8.6.2 Interface Proxy-ARP

ip address

ip address <ip-address>
no ip address

Sets IP address for this interface.
 The no form of the command sets the ip address to 0.0.0.0.

Syntax Description	ip-address	IP address.
Default	0.0.0.0	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.3000	First version
	3.3.4100	Removed “prefix” parameter
	3.3.4302	Added note
	3.3.4402	Added HA note
Role	admin	
Example	switch (config interface proxy-arp 1) # ip-address 10.10.10.10 switch (config interface proxy-arp 1) #	
Related Commands		
Note	• The no form of the command does not affect the mask • The subnets of different Proxy-ARP interfaces must not overlap or conflict with one another • If the IP address is set to 0.0.0.0 traffic is not forwarded by the Proxy-ARP. The IP address must be set before Ethernet-InfiniBand or InfiniBand-Ethernet traffic is passed. • When Proxy-ARP HA is enabled, the command is accessible from VIP. When Proxy-ARP HA is disabled, the command is accessible from the gateway management interface.	

ip netmask

ip netmask <mask>
no ip netmask

Sets mask length for the IP address.
The no form of the command resets this parameter to its default value.

Syntax Description	mask	The IP mask. Input format can be x.y.z.w (e.g. 255.255.255.0), or mask length (range 0-32) may be provided (e.g. /24).
Default	0	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4100	First version
	3.3.4402	Added HA note
Role	admin	
Example	switch (config interface proxy-arp 1) # ip netmask /24 switch (config interface proxy-arp 1) #	
Related Commands		
Note	• Make sure the IP address is configured using the command “ip address” • This command only configures the subnet mask • If Proxy-ARP VIP is used this command is applicable via the VIP only.	

ip vlan

ip vlan <vlan-id>
no ip vlan

Sets the VLAN number.
 The no form of the command deletes the VLAN.

Syntax Description	vlan-id	12-bit VLAN ID. The range: [1-4094].
Default	No VLAN is configured	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4000 3.3.4100 Added ability to edit VLAN ID. 3.3.4302 Updated Note field 3.3.4402 Added HA note	
Role	admin	
Example	<pre>switch (config interface proxy-arp 1) # ip vlan 10 switch (config interface proxy-arp 1) #</pre>	
Related Commands		
Note	- VLAN must be configured on the Ethernet interface to enable gateway functionality. - Each Proxy-ARP interface must have a unique VLAN/PKEY pair. Trying to map the same VLAN/PKEY to a different Proxy-ARP interface results in an error message. - If Proxy-ARP VIP is used this command is applicable via the VIP only.	

ip pkey

ip pkey <pkey>
no ip pkey

Sets the PKEY number (InfiniBand partition).
The no form of the command deletes the PKEY.

Syntax Description	pkey	The range is [0x1 – 0x7fff]. Partition value can be specified as a base 10 number or a hexa value (with “0x” prefix).
Default	No PKEY is configured	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4000	
	3.3.4100	Added ability to edit PKEY number.
	3.3.4302	Updated Note field
	3.3.4402	Added HA note
Role	admin	
Example	switch (config interface proxy-arp 1) # ip pkey 0x7fff switch (config interface proxy-arp 1) #	
Related Commands		
Note	- The PKEY value is always reported in hexadecimal values via the command “show interfaces proxy-arp”. - Each Proxy-ARP interface must have a unique VLAN/PKEY pair. Trying to map the same VLAN/PKEY to a different Proxy-ARP interface results in an error message. - If Proxy-ARP VIP is used this command is applicable via the VIP only.	

ip route

ip route <ip-address/prefix> <next-hop-ip-address>
no ip route [<ip-address/prefix>]

Adds a static route via this interface.
 The no form of the command deletes a specified IP route or all IP routes for this interface.

Syntax Description	ip-address/prefix	IP address and network prefix (for example, 10.10.10.0/24)
	next-hop-ip-address	The next hop IP address for this route – the default router IP address
Default	N/A	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4000	First version
	3.3.4302	Added static routes note
	3.3.4402	Added HA note
	3.4.0000	Updated no form
Role	admin	
Example	switch (config interface proxy-arp 1) # ip route 10.10.10.0/24 10.10.10.254	
Related Commands		
Note	• 160 static routes can be configured altogether allocated for the 32 Proxy-ARP interfaces. • If Proxy ARP VIP is used, this command is applicable via the VIP only.	

ip route default

ip route default <next-hop-ip-address>
no ip route default <next-hop-ip-address>

Adds a default static route via this interface.
The no form of the command deletes the default static route.

Syntax Description	next-hop-ip-address	The next hop IP address for this route – the default router IP address
Default	N/A	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.0000	First version
	3.3.4402	Added HA note
Role	admin	
Example	switch (config interface proxy-arp 1) # ip route default 10.10.10.254 switch (config interface proxy-arp 1) #	
Related Commands		
Note	If Proxy ARP VIP is used, this command is applicable via the VIP only.	

description

description <string>
no description

Sets a user-defined string for the description.
 The no form of the command sets the parameter to default.

Syntax Description	string	40B string of user-defined description for the interface
Default	"" (empty string)	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.0000	First version
	3.3.4402	Added HA note
	3.4.0000	Updated note
Role	admin	
Example	<pre>switch (config interface proxy-arp 1) # description my-description switch (config interface proxy-arp 1) #</pre>	
Related Commands		
Note	If Proxy ARP VIP is used, this command is applicable via the VIP only.	

shutdown

shutdown
no shutdown

Disables the interface.
The no form of the command enables the interface.

Syntax Description	N/A	
Default	Disabled	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.0000	First version
	3.3.4402	Added HA note
Role	admin	
Example	switch (config interface proxy-arp 1) # no shutdown switch (config interface proxy-arp 1) #	
Related Commands		
Note	If Proxy-ARP VIP is used this, command is applicable via VIP only.	

mtu

mtu <value>

no mtu

Sets the MTU of the interface.

The no form of the command sets the MTU to default.

Syntax Description	value	The MTU value. Range: 1500-4092 bytes.
Default	1500	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4000	First version
	3.3.4402	Added HA note
Role	admin	
Example	switch (config interface proxy-arp 1) # mtu 2044 switch (config interface proxy-arp 1) #	
Related Commands		
Note	If Proxy-ARP VIP is used, this command is applicable via VIP only.	

counters

counters
no counters

Enables counters on this interface.
The no form of the command disables counters on this interface.

Syntax Description	N/A
Default	Counters are disabled.
Configuration Mode	Config Interface Proxy-ARP
History	3.3.4000 First version 3.3.4302 Added note 3.3.4402 Added HA note
Role	admin
Example	<pre>switch (config interface proxy-arp 1) # counters switch (config interface proxy-arp 1) #</pre>
Related Commands	
Note	- The system is limited in the number of L3 counters, not all interfaces may have counters enabled - To see counters on the interface you need first to enable counters by this command - Counters are supported on up to 12 Proxy-ARP interfaces simultaneously. Enabling counters on more than 12 interfaces is allowed, but results in warnings in the log and causes uncertain behavior. - If Proxy ARP VIP is used, this command is applicable via VIP only.

clear counters

clear counters

Clears interface counters.

Syntax Description	N/A	
Default	N/A	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4000	First version
	3.3.4402	Added HA note
Role	admin	
Example	switch (config interface proxy-arp 1) # clear counters switch (config interface proxy-arp 1) #	
Related Commands		
Note	• If the command is running from the VIP it cleans the counters from all Proxy-ARP group members. • If the command is running from the gateway node IP it clears only the specific member's counters.	

ha table-update-interval

ha table-update-interval <seconds>
no ha table-update-interval

Sets the time the standby table is considered valid.
The no form of the command resets the parameter to its default value.

Syntax Description	seconds	Time in seconds. Range: 1-10000.
Default	120	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4402	
Role	admin	
Example	switch (config interface proxy-arp 1) # ha table-update-interval 120 switch (config interface proxy-arp 1) #	
Related Commands		
Note	If Proxy-ARP VIP is used, this command is applicable via VIP only.	

ha table-fast-learn-time

ha table-fast-learn-time <seconds>
no ha table-fast-learn-time

Sets the time a new master (after the previous master failed) waits before forwarding the load balancing vector to the other (non-master) nodes in the system.

The no form of the command resets the parameter to its default value.

Syntax Description	seconds	Time in seconds. Range: 1-10000.
Default	1	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4402	
Role	admin	
Example	<pre>switch (config interface proxy-arp 1) # ha table-fast-learn-time 1 switch (config interface proxy-arp 1) #</pre>	
Related Commands	ha master-election-learn-interval	
Note	- This command is only available on the Proxy-ARP HA group master and via VIP. - The value of table-fast-learn-time must be smaller than the value of table-learn-time.	

ha master-election-learn-interval

ha master-election-learn-interval <seconds>
no ha master-election-learn-interval

Sets the time a new Proxy-ARP group master stays in learning mode waiting for an advertisement to be received.

The no form of the command resets the parameter to its default value.

Syntax Description	seconds	Time in seconds. Range: 1-10000.
Default	30	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4402	
Role	admin	
Example	<pre>switch (config interface proxy-arp 1) # ha master-election-learn-interval 30 switch (config interface proxy-arp 1) #</pre>	
Related Commands	ha table-fast-learn-time	
Note	- Once this timer expires, this node becomes the Proxy-ARP group master if no master advertisements are received. - If Proxy-ARP VIP is used, this command is applicable via VIP only. - The value of master-election-learn-interval must be greater than the value of advertisement-interval.	

ha advertisement-interval

ha advertisement-interval <seconds>
no ha advertisement-interval

Sets the time between keepalive messages (sent by the standby node) and master advertisement messages (sent by the master node).
 The no form of the command resets the parameter to its default value.

Syntax Description	seconds	Time in seconds. Range: 1-10000.
Default	1	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4402	
Role	admin	
Example	switch (config interface proxy-arp 1) # ha advertisement-interval 1 switch (config interface proxy-arp 1) #	
Related Commands	ha table-learn-time	
Note	• If Proxy-ARP VIP is used, this command is applicable via VIP only. • The value of advertisement-interval must be smaller than the value of master-election-learn-interval.	

ha table-learn-time

ha table-learn-time <seconds>
no ha table-learn-time

Sets the time a new master (when no previous master is known) waits before forwarding the load balancing vector to the other (non-master) nodes in the system.

The no form of the command resets the parameter to its default value.

Syntax Description	seconds	Time in seconds. Range: 1-10000.
Default	30	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4402	
Role	admin	
Example	switch (config interface proxy-arp 1) # ha table-learn-time 30 switch (config interface proxy-arp 1) #	
Related Commands	ha advertisement-interval	
Note	- This command is only available on the Proxy-ARP HA group master and via VIP. - The value of table-learn-time must be greater than the value of table-fast-learn-time.	

ha keep-alive-loss-threshold

ha keep-alive-loss-threshold <value>
no ha keep-alive-loss-threshold

Sets the number of lost keep alive messages in the keep alive loss interval that cause an alarm to be generated.
 The no form of the command resets the parameter to its default value.

Syntax Description	value	Time in seconds. Range: 1-10000.
Default	10	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4402	
	3.3.4500	Updated Note section
Role	admin	
Example	switch (config interface proxy-arp 1) # ha keep-alive-loss-threshold 10 switch (config interface proxy-arp 1) #	
Related Commands	ha keep-alive-loss-interval	
Note	- If Proxy-ARP VIP is used, this command is applicable via VIP only. - The alarm generated is a WARNING message in the log.	

ha keep-alive-loss-interval

ha keep-alive-loss-interval <seconds>
no ha keep-alive-loss-interval

Sets the time interval during which lost keepalive messages (more than what is configured in “ha keep-alive-loss-threshold”) trigger an alarm.
The no form of the command resets the parameter to its default value.

Syntax Description	seconds	Time in seconds. Range: 1-10000.
Default	120	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4402	
	3.3.4500	Updated Note section
Role	admin	
Example	switch (config interface proxy-arp 1) # ha keep-alive-loss-interval 120 switch (config interface proxy-arp 1) #	
Related Commands	ha keep-alive-loss-threshold	
Note	• If Proxy-ARP VIP is used, this command is applicable via VIP only. • The alarm generated is a WARNING message in the log.	

ha host-list-diff-update-interval

ha host-list-diff-update-interval <seconds>
no ha host-list-diff-update-interval

Sets the time to wait after receiving a new host before sending a host list update message to the master.

The no form of the command resets the parameter to its default value.

Syntax Description	seconds	Time in seconds. Range: 1-10000.
Default	3	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4402	
Role	admin	
Example	<pre>switch (config interface proxy-arp 1) # ha host-list-diff-update-interval 3 switch (config interface proxy-arp 1) #</pre>	
Related Commands	ha host-list-update-interval	
Note	• If Proxy-ARP VIP is used, this command is applicable via VIP only. • The value of host-list-diff-update-interval must be smaller than the value of host-list-update-interval.	

ha host-list-update-interval

ha host-list-update-interval <seconds>
no ha host-list-update-interval

Sets the number of seconds non-master nodes send their full host list to the master.

The no form of the command resets the parameter to its default value.

Syntax Description	seconds	Time in seconds. Range: 1-10000.
Default	300	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4402	
Role	admin	
Example	<pre>switch (config interface proxy-arp 1) # ha host-list-diff-update-interval 300 switch (config interface proxy-arp 1) #</pre>	
Related Commands	ha host-list-diff-update-interval	
Note	• If Proxy-ARP VIP is used, this command is applicable via VIP only. • The value of host-list-update-interval must be greater than the value of host-list-diff-update-interval. • It is recommended not to set host-list-update-interval to be less than 2 minutes when working under heavy traffic.	

ha member description

ha member <hostname> description <text>
no ha member <hostname> description

Sets a description per member in the Proxy-ARP group.
 The no form of the command resets the parameter to its default value.

Syntax Description	hostname	The member's hostname
	text	User defined description. String range is up to 40 bytes.
Default	""	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4402	
Role	admin	
Example	<pre>switch (config interface proxy-arp 1) # ha member my-member-hostname description my-member-description switch (config interface proxy-arp 1) #</pre>	
Related Commands		
Note	If Proxy-ARP VIP is used, this command is applicable via VIP only.	

ha member ip address

ha member <hostname> ip address <ip-address>
no ha member <hostname> ip address <ip-address>

Assigns an IP address to access the instance on this HA member.
The no form of the command resets the parameter to its default value.

Syntax Description	hostname	The member's hostname
	ip-address	IPv4 address in a dotted format
Default	0.0.0.0	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4402	
	3.3.4500	Added a note
Role	admin	
Example	switch (config interface proxy-arp 1) # ha member my-member-hostname ip address 10.10.10.10 switch (config interface proxy-arp 1) #	
Related Commands		
Note	• If Proxy-ARP VIP is used, this command is applicable via VIP only. • The network mask is configured via the command “interface proxy-arp <pra-id> ip masklen <value>” • The IP address configured must be that of the management interface of the system.	

ha member priority

ha member <hostname> priority <value>
no ha member <hostname> priority

Sets the master election priority for member when servicing the Proxy-ARP interface.

The no form of the command resets the parameter to its default value.

Syntax Description	hostname	The member's hostname
	value	The priority value. Range is 1-255.
Default	100	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4402	
Role	admin	
Example	switch (config interface proxy-arp 1) # ha member my-member-hostname priority 100 switch (config interface proxy-arp 1) #	
Related Commands		
Note	• If Proxy-ARP VIP is used, this command is applicable via VIP only. • Priority is calculated based on the configured priority and the GUID. The active member with the highest priority becomes master the next time master election takes place.	

ha member clear counters

ha member <hostname> clear counters

Clears a specific member's interface counters.

Syntax Description	hostname	The member's hostname
Default	N/A	
Configuration Mode	Config Interface Proxy-ARP	
History	3.3.4402	
Role	admin	
Example	<pre>switch (config interface proxy-arp 1) # ha member my-member-hostname clear counters switch (config interface proxy-arp 1) #</pre>	
Related Commands		
Note	If Proxy-ARP VIP is used, this command is applicable via VIP only.	

ip multicast

ip multicast <ipv4-address>
no ip multicast [<ipv4-address>]

Adds a multicast group to a Proxy-ARP instance.
 The no form of the command deletes a specified multicast group from a Proxy-ARP instance or all multicast groups that Proxy-ARP instance belongs to.

Syntax Description	ipv4-address	Specifies a multicast group. Must be a valid IPv4 address in the MC range (224.0.0.0 – 239.255.255.255 – however, the address 239.0.0.77 cannot be given, since this address is reserved for GLBP).
Default	N/A	
Configuration Mode	Config Interface Proxy-ARP	
History	3.4.0000	
Role	admin	
Example	switch (config interface proxy-arp 1) # ip multicast [<ipv4-address>]	
Related Commands		
Note	• The command is available when the platform is in standalone mode only, and only in Proxy-ARP MC mode. • The groups 224.0.0.1 (all-nodes group) and 224.0.0.2 (all-routers group) can be configured but will actually not have any effect on the configuration of these routes in the hardware. These two groups are configured automatically by the application, regardless of user configuration.	

ip multicast filter

ip multicast filter <ipv4-address> <ipv4-netmask>
no ip multicast filter [<ipv4-address>] [<ipv4-netmask>]

Adds multicast group based on netmask to a Proxy-ARP instance filtering table.

The no form of the command deletes multicast group.

Syntax Description	ipv4-address	The multicast group IPv4 address
	ipv4-netmask	The multicast group IPv4 address
Default	N/A	
Configuration Mode	Config Interface Proxy-ARP	
History	3.4.3100	
Role	admin	
Example	switch (config interface proxy-arp 1) # ip multicast filter 10.10.10.10 255.255.0.0	
Related Commands		
Note	• The IPv4 address must be a valid IPv4 address in the multicast range (Class D: 224.0.0.0 – 239.255.255.255) • The command is available in single-box mode, or in HA mode from the VIP • IPv4 netmask can be set in two ways: 255.255.255.255 or /32 • A valid netmask must be set • If a static multicast group is configured within the MC filter range it is not blocked • If in the no form of the command no IPv4 address is provided, all multi-cast groups that the Proxy-ARP ID belongs to are deleted	

8.6.3 Show

show ip proxy-arp

show ip proxy-arp

Displays the admin state of the Proxy-ARP feature.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.0000
	3.3.4402 Added note.
Role	admin; monitor
Example	switch (config) # show ip proxy-arp Proxy-arp: enabled switch (config) #
Related Commands	
Note	This command is applicable from the gateway node IP only.

show ip route

show ip route [interface proxy-arp <pra-id>]

Displays routing table.

Syntax Description	interface proxy-arp <pra-id>	Displays IP route table for a specific Proxy-ARP interface
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.4000 3.3.4402	Added note.
Role	admin; monitor	
Example	<pre>switch (config) # show ip route interface proxy-arp 1 Destination Mask Gateway Interface Source Distance/Metric 22.22.0.0 255.255.0.0 0.0.0.0 proxy-arp 1 kernel 0/0 default 0.0.0.0 22.22.0.174 proxy-arp 1 static 0/0 switch (config) #</pre>	
Related Commands		
Note	This command is applicable from the gateway node IP only.	

show ip arp

show ip arp [interface proxy-arp [<instance>] ha designated-member]

Displays ARP table.

Syntax Description	interface proxy-arp [<instance>]	Filters the table according to a specific interface (i.e. proxy-arp, mgmt0)												
Default	N/A													
Configuration Mode	Any Command Mode													
History	3.3.3000 3.3.4402	Added note.												
Role	admin; monitor													
Example	switch (config) # show ip arp interface proxy arp 1 Total number of entries: 2 <table><thead><tr><th>Address</th><th>Type</th><th>Hardware Address</th><th>Interface</th></tr></thead><tbody><tr><td>192.168.12.25</td><td>Dynamic ETH</td><td>00:02:C9:45:45:51</td><td>proxy-arp 1</td></tr><tr><td>192.168.12.15</td><td>Infiniband</td><td>00:02:C9:02:00:24:CB:4</td><td>proxy-arp 1</td></tr></tbody></table> switch (config) #		Address	Type	Hardware Address	Interface	192.168.12.25	Dynamic ETH	00:02:C9:45:45:51	proxy-arp 1	192.168.12.15	Infiniband	00:02:C9:02:00:24:CB:4	proxy-arp 1
Address	Type	Hardware Address	Interface											
192.168.12.25	Dynamic ETH	00:02:C9:45:45:51	proxy-arp 1											
192.168.12.15	Infiniband	00:02:C9:02:00:24:CB:4	proxy-arp 1											
Related Commands														
Note	- This command is applicable from the gateway node IP only. - If Proxy-ARP HA is enabled. the parameter “ha designated-member” makes the ARP table include only the entries designated by this specific node.													

show proxy-arp dhcp

show proxy-arp dhcp

Displays Proxy-ARP DHCP mode.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.4.2100
Role	admin
Example	switch (config) # show proxy-arp dhcp Proxy-arp dhcp: linux switch (config) #
Related Commands	
Note	

show interfaces proxy-arp

show interfaces proxy-arp <id> [configured]

Displays interface configuration and status.

Syntax Description	id	A Proxy-ARP interface number
	configured	Displays information only about a Proxy-ARP interface that has been configured by the user
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.0000	
	3.3.4402	Added note.
Role	admin	
Example	<pre>switch (config) # show interfaces proxy-arp 1 configured Proxy-arp 1 Admin state: Enabled Internet Address: 108.10.24.110/16 Description: N/A MTU: 1500 Counters: Enabled Member interfaces: vlan 3374, pkey 0x7fff switch (config) #</pre>	
Related Commands		
Note	This command is applicable from the gateway node IP only.	

show interfaces proxy-arp brief

show interfaces proxy-arp brief

Displays the status of interfaces.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.4000 3.3.4402 Added note.
Role	admin
Example	<pre>switch (config)# show interfaces proxy-arp brief Interface Description State Bridged interfaces ----- proxy-arp 1 descr1 Up vlan 23, pkey 0x23 proxy-arp 2 descr2 Down vlan 24, pkey 0x24 switch (config)#</pre>
Related Commands	
Note	This command is applicable from the management node IP only.

show sma port

show sma port <number>

Displays vTCA port configuration.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.0000
	3.3.4402 Added note.
Role	admin
Example	switch (config) # show sma port 1 Enabled switch (config) #
Related Commands	
Note	This command is applicable from the gateway node IP only.

show proxy-arp ha

show proxy-arp ha

Displays the configuration and status of all nodes in the HA Proxy-ARP group.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.4402
Role	admin
Example	<pre>switch (config) # show proxy-arp ha Load balancing: ib-base-ip Number of Proxy-Arp interfaces: 32 Proxy Arp VIP: ===== Proxy-arp group name: my-pra-group HA VIP address: 192.168.11.55/24 Active nodes: 4 Hostname State IP Address ----- gateway1 master 10.10.10.11 gateway2 standby 10.10.10.12 gateway3 standby 10.10.10.13 gateway4 standby 10.10.10.14 switch (config) #</pre>
Related Commands	
Note	• If Proxy ARP VIP is used, this command is applicable via the VIP only. • Up to 16 active nodes from the same CPU type are supported. • The IP addresses displayed is the management IP of the gateways.

show interfaces proxy-arp ha

show interfaces proxy-arp <pra-id> ha

Displays the Proxy-ARP interface configuration and a summary of per-member configuration and status.

Syntax Description	pra-id	Proxy-ARP number
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.4402	
Role	admin	
Example	<pre>switch (config) # show interfaces proxy-arp 1 ha Proxy-arp 1 Keep Alive loss threshold: 10 Keep Alive loss interval: 120 seconds Load balancing algorithm: ib-base-ip IP masklen: 24 Admin state: Enabled MTU: 1500 Counters: Disabled Bridged interfaces: vlan 10, pkey 0x7fff Number of members: 3 Hostname Description Admin State LB State Operational State IP Priority ----- hostname1 host1-dscr Enabled Active Up 192.168.12.30 100 hostname2 host2-dscr Enabled Standby Up 192.168.12.33 100 hostname3 host3-dscr Enabled Down Down 192.168.12.31 100 switch (config) #</pre>	
Related Commands		
Note	This command is applicable via the VIP only.	

show interfaces proxy-arp ha ip arp

show interfaces proxy-arp <pra-id> ha ip arp [designated]

Displays the current ARP table of all members of this Proxy-ARP instance.

Syntax Description	pra-id	Proxy-ARP ID																																								
	designated	Displays the current designated ARP table for all members of this Proxy-ARP instance																																								
Default	N/A																																									
Configuration Mode	Any Command Mode																																									
History	3.3.4402																																									
Role	admin																																									
Example	<pre>switch (config) # show interfaces proxy-arp 1 ha ip arp</pre> Proxy-arp 1 HA arp table: Total number of entries: 9 <table><thead><tr><th>Address</th><th>Type</th><th>Hardware Address</th><th>Members</th></tr></thead><tbody><tr><td>192.167.5.2</td><td>Dynamic ETH</td><td>00:02:C9:43:1C:D1</td><td>r-mgtswd-125</td></tr><tr><td>192.167.5.3</td><td>Dynamic IB</td><td>00:14:05:00:00:00:00:80</td><td>r-mgtswd-125, r-mgtswd-126</td></tr><tr><td>192.167.5.4</td><td>Dynamic IB</td><td>00:14:05:00:00:00:00:80</td><td>r-mgtswd-125</td></tr><tr><td>192.167.5.5</td><td>Dynamic IB</td><td>00:14:05:00:00:00:00:80</td><td>r-mgtswd-125</td></tr><tr><td>192.167.5.6</td><td>Dynamic IB</td><td>00:14:05:00:00:00:00:80</td><td>r-mgtswd-125</td></tr><tr><td>192.167.5.7</td><td>Dynamic IB</td><td>00:14:05:00:00:00:00:80</td><td>r-mgtswd-125</td></tr><tr><td>192.167.5.8</td><td>Dynamic IB</td><td>00:14:05:00:00:00:00:80</td><td>r-mgtswd-125</td></tr><tr><td>192.167.5.9</td><td>Dynamic IB</td><td>00:14:05:00:00:00:00:80</td><td>r-mgtswd-125</td></tr><tr><td>192.167.5.10</td><td>Dynamic IB</td><td>00:14:05:00:00:00:00:80</td><td>r-mgtswd-125</td></tr></tbody></table> <pre>switch (config) #</pre>		Address	Type	Hardware Address	Members	192.167.5.2	Dynamic ETH	00:02:C9:43:1C:D1	r-mgtswd-125	192.167.5.3	Dynamic IB	00:14:05:00:00:00:00:80	r-mgtswd-125, r-mgtswd-126	192.167.5.4	Dynamic IB	00:14:05:00:00:00:00:80	r-mgtswd-125	192.167.5.5	Dynamic IB	00:14:05:00:00:00:00:80	r-mgtswd-125	192.167.5.6	Dynamic IB	00:14:05:00:00:00:00:80	r-mgtswd-125	192.167.5.7	Dynamic IB	00:14:05:00:00:00:00:80	r-mgtswd-125	192.167.5.8	Dynamic IB	00:14:05:00:00:00:00:80	r-mgtswd-125	192.167.5.9	Dynamic IB	00:14:05:00:00:00:00:80	r-mgtswd-125	192.167.5.10	Dynamic IB	00:14:05:00:00:00:00:80	r-mgtswd-125
Address	Type	Hardware Address	Members																																							
192.167.5.2	Dynamic ETH	00:02:C9:43:1C:D1	r-mgtswd-125																																							
192.167.5.3	Dynamic IB	00:14:05:00:00:00:00:80	r-mgtswd-125, r-mgtswd-126																																							
192.167.5.4	Dynamic IB	00:14:05:00:00:00:00:80	r-mgtswd-125																																							
192.167.5.5	Dynamic IB	00:14:05:00:00:00:00:80	r-mgtswd-125																																							
192.167.5.6	Dynamic IB	00:14:05:00:00:00:00:80	r-mgtswd-125																																							
192.167.5.7	Dynamic IB	00:14:05:00:00:00:00:80	r-mgtswd-125																																							
192.167.5.8	Dynamic IB	00:14:05:00:00:00:00:80	r-mgtswd-125																																							
192.167.5.9	Dynamic IB	00:14:05:00:00:00:00:80	r-mgtswd-125																																							
192.167.5.10	Dynamic IB	00:14:05:00:00:00:00:80	r-mgtswd-125																																							
Related Commands																																										
Note	• If Proxy ARP VIP is used, this command is applicable via the VIP only. • Aged-out ARP entries may also appear.																																									

show interfaces proxy-arp ha designated-member

show interfaces proxy-arp <pra-id> ha designated-member <Ethernet-IP> <InfiniBand-IP>

Calculates which member of this Proxy-ARP instance handles the traffic between the given Ethernet IP and the given InfiniBand IP.

Syntax Description	pra-id	Proxy-ARP ID
	IP-Address	Host IP connected to the Ethernet network
	InfiniBand-IP	Host IP connected to the InfiniBand network
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.4402	
Role	admin	
Example	<pre>switch (config) # show interfaces proxy-arp 1 ha designated-member 10.10.10.10 10.10.10.11 Member: <calculated hostname that will pass the traffic> switch (config) #</pre>	
Related Commands		
Note	This command is applicable via the VIP only.	

show interfaces proxy-arp ha host-list

show interfaces proxy-arp <pra-id> ha host-list designated-member

Displays designated member for passing Ethernet-IB traffic

Syntax Description	pra-id				Proxy-ARP ID												
Default	N/A																
Configuration Mode	Any Command Mode																
History	3.3.4402																
Role	admin																
Example	<pre>switch (config) # show interfaces proxy-arp 7 ha host-list designated-member</pre> Proxy-arp 7 HA host list: Total number of entries: 2 <table><thead><tr><th>Address</th><th>Type</th><th>Hardware Address</th><th>Members</th></tr></thead><tbody><tr><td>7.120.10.55</td><td>Dynamic ETH</td><td>00:02:C9:45:5B:40</td><td>Piranha-GW-242</td></tr><tr><td>7.120.10.54</td><td>Dynamic IB</td><td>00:02:C9:03:00:E6:E5:41</td><td>Piranha-GW-242</td></tr></tbody></table> <pre>switch (config) #</pre>					Address	Type	Hardware Address	Members	7.120.10.55	Dynamic ETH	00:02:C9:45:5B:40	Piranha-GW-242	7.120.10.54	Dynamic IB	00:02:C9:03:00:E6:E5:41	Piranha-GW-242
Address	Type	Hardware Address	Members														
7.120.10.55	Dynamic ETH	00:02:C9:45:5B:40	Piranha-GW-242														
7.120.10.54	Dynamic IB	00:02:C9:03:00:E6:E5:41	Piranha-GW-242														
Related Commands																	
Note	This command is applicable via the VIP only.																

show interfaces proxy-arp ha detail

show interfaces proxy-arp <pra-id> ha detail

Displays the Proxy ARP interface detail configuration and provides a summary of per-member configuration and status.

Syntax Description	pra-id	Proxy-ARP ID
	Ethernet-IP	Host IP connected to the Ethernet network
	InfiniBand-IP	Host IP connected to the InfiniBand network
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.4402	
Role	admin	
Example	switch (config) # show interfaces proxy-arp 1 ha detail	

```
Proxy-arp 1
  Table update interval: 120 seconds
  PRA table fast learn time: 1 second
  Master election learning interval: 30 seconds
  Advertisement interval: 1 second
  PRA table learn time: 30 seconds
  Keep Alive loss threshold: 10
  Keep Alive loss interval: 120 seconds
  Host list differential update interval: 3 seconds
  Host list update interval: 300 seconds
  Load balancing algorithm: ib-base-ip
  IP masklen: 24
  Admin state: Enabled
  MTU: 1500
  Counters: Disabled
  Bridged interfaces: vlan 2, pkey 0x3
```

Hostname	Description	Admin State	LB State	Operational	State IP	Priority
hostname1*	host1-dscr	Enabled	Active	Up	192.168.12.30	100
hostname2	host2-dscr	Enabled	Standby	Up	192.168.12.33	100
hostname3	host3-dscr	Enabled	Down	Down	192.168.12.31	100

```
switch (config) #
```

Related Commands

Note This command is applicable via the VIP only.

show interfaces proxy-arp ha member

show interfaces proxy-arp <pra-id> ha member <hostname>

Displays member interface configuration and status.

Syntax Description	pra-id	Interface number
	hostname	Member's hostname
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.4402	
Role	admin	
Example	<pre>switch (config) # show interfaces proxy-arp 1 ha member my-member Proxy-arp 1 Admin state: Enabled Operational state: Up GUID: 00:02:C9:03:00:66:08:63 Internet Address: 10.10.10.10/24 Broadcast Address: 10.10.10.255 Description: N/A MTU: 1500 Counters: Disabled Bridged interfaces: vlan 10, pkey 0x7fff switch (config) #</pre>	
Related Commands		
Note	This command is applicable via the VIP only.	

show proxy-arp mode

show proxy-arp mode

Displays the Proxy-ARP mode.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.4.0000
	3.6.3004 Updated Example
Role	admin
Example	switch (config) # show proxy-arp mode Proxy-arp mode: multicast switch (config) #
Related Commands	
Note	Resources may be found using the command “show system resource table”

show ip multicast interface proxy-arp

show ip multicast interface proxy-arp [<if-number> | <if-number> debug]

Displays multicast table for all Proxy-ARP interfaces.

Syntax Description	if-number	Specifies interface number
	debug	Displays Proxy-ARP instance multicast information including debug information
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.4.0000	
	3.4.1100	Updated Example
Role	admin	
Example	<pre>switch (config) # show ip multicast interface proxy-arp Proxy-arp multicast list: Total number of entries: 2 Address Interface Source ----- 239.0.0.77 proxy-arp 2 Dynamic 255.255.255.255 proxy-arp 2 Dynamic switch (config) #</pre>	
Related Commands		
Note	This command is applicable from the gateway node IP only.	

show ip multicast interface proxy-arp count

show ip multicast interface proxy-arp count

Displays number of multicast groups for all Proxy-ARP interfaces.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.4.1100
Role	admin
Example	switch (config) # show ip multicast interface proxy-arp count Proxy-arp multicast count: 20 (static: 12, dynamic: 8) switch (config) #
Related Commands	
Note	This command is applicable from the gateway node IP only.

show interfaces proxy-arp ha multicast-list

show interfaces proxy-arp <pra-id> ha multicast-list

Displays multicast table per Proxy-ARP interface, for all cluster members.

Syntax Description	pra-id	Proxy-ARP number
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.4.1100	
Role	admin	
Example	<pre>switch (config) # show interfaces proxy-arp 70 ha multicast-list Proxy-arp 70 HA multicast list: Total number of entries: 10 Address Source Members ----- 224.0.0.1 Dynamic x86-GW-181 224.0.0.2 Dynamic x86-GW-181 228.0.0.0 Static x86-GW-181 228.0.0.1 Static x86-GW-181 228.0.0.2 Static x86-GW-181 228.0.0.3 Static x86-GW-181 228.0.0.4 Static x86-GW-181 228.0.0.5 Static x86-GW-181 239.0.0.77 Dynamic x86-GW-181 255.255.255.255 Dynamic x86-GW-181 switch (config) #</pre>	
Related Commands		
Note	This command is applicable from the VIP only.	

show ip multicast filter interface proxy-arp

show ip multicast filter interface proxy-arp [<pra-id>]

Displays all filtered multicast routes configured to all or a specific Proxy-ARP on this gateway node.

Syntax Description	pra-id	Proxy-ARP number
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.4.3100	
Role	admin	
Example	<pre>switch (config) # show ip multicast filter interface proxy-arp 1 Address Mask Interface ----- 224.0.0.5 255.255.255.255 proxy-arp 1 224.0.0.6 255.255.255.255 proxy-arp 2 229.0.0.0 255.255.0.0 proxy-arp 2</pre>	
Related Commands		
Note	- This command displays MC filter route configuration as it appears in the kernel. To view all configured routes, run the command “show running-config”. - This command is only available from the gateway node IP	

show interfaces proxy-arp ha multicast-filter-list

show interfaces proxy-arp <pra-id> ha multicast-filter-list

Displays all active multicast filter routes configured to a specific Proxy-ARP on all nodes part of this Gateway group.

Syntax Description	pra-id	Proxy-ARP number																																																						
Default	N/A																																																							
Configuration Mode	Any Command Mode																																																							
History	3.4.3100																																																							
Role	admin																																																							
Example	<pre>switch (config) # show interfaces proxy-arp 66 ha multicast-filter-list</pre> Total number of entries: 101 <table> <thead> <tr> <th>Address</th><th>Mask</th><th>Interface</th></tr> </thead> <tbody> <tr><td>224.110.249.32</td><td>255.255.255.224</td><td>proxy-arp 66</td></tr> <tr><td>224.128.0.0</td><td>255.240.0.0</td><td>proxy-arp 66</td></tr> <tr><td>224.140.64.0</td><td>255.255.192.0</td><td>proxy-arp 66</td></tr> <tr><td>224.180.0.0</td><td>255.255.0.0</td><td>proxy-arp 66</td></tr> <tr><td>224.227.149.0</td><td>255.255.255.128</td><td>proxy-arp 66</td></tr> <tr><td>224.251.0.0</td><td>255.255.224.0</td><td>proxy-arp 66</td></tr> <tr><td>225.69.118.64</td><td>255.255.255.224</td><td>proxy-arp 66</td></tr> <tr><td>225.97.128.0</td><td>255.255.128.0</td><td>proxy-arp 66</td></tr> <tr><td>225.128.0.0</td><td>255.128.0.0</td><td>proxy-arp 66</td></tr> <tr><td>225.134.0.0</td><td>255.255.192.0</td><td>proxy-arp 66</td></tr> <tr><td>225.234.0.0</td><td>255.255.0.0</td><td>proxy-arp 66</td></tr> <tr><td>225.234.207.128</td><td>255.255.255.128</td><td>proxy-arp 66</td></tr> <tr><td>226.0.0.0</td><td>255.128.0.0</td><td>proxy-arp 66</td></tr> <tr><td>226.14.9.20</td><td>255.255.255.254</td><td>proxy-arp 66</td></tr> <tr><td>226.32.0.0</td><td>255.254.0.0</td><td>proxy-arp 66</td></tr> <tr><td>226.108.76.0</td><td>255.255.254.0</td><td>proxy-arp 66</td></tr> <tr><td>...</td><td></td><td></td></tr> </tbody> </table>		Address	Mask	Interface	224.110.249.32	255.255.255.224	proxy-arp 66	224.128.0.0	255.240.0.0	proxy-arp 66	224.140.64.0	255.255.192.0	proxy-arp 66	224.180.0.0	255.255.0.0	proxy-arp 66	224.227.149.0	255.255.255.128	proxy-arp 66	224.251.0.0	255.255.224.0	proxy-arp 66	225.69.118.64	255.255.255.224	proxy-arp 66	225.97.128.0	255.255.128.0	proxy-arp 66	225.128.0.0	255.128.0.0	proxy-arp 66	225.134.0.0	255.255.192.0	proxy-arp 66	225.234.0.0	255.255.0.0	proxy-arp 66	225.234.207.128	255.255.255.128	proxy-arp 66	226.0.0.0	255.128.0.0	proxy-arp 66	226.14.9.20	255.255.255.254	proxy-arp 66	226.32.0.0	255.254.0.0	proxy-arp 66	226.108.76.0	255.255.254.0	proxy-arp 66	...		
Address	Mask	Interface																																																						
224.110.249.32	255.255.255.224	proxy-arp 66																																																						
224.128.0.0	255.240.0.0	proxy-arp 66																																																						
224.140.64.0	255.255.192.0	proxy-arp 66																																																						
224.180.0.0	255.255.0.0	proxy-arp 66																																																						
224.227.149.0	255.255.255.128	proxy-arp 66																																																						
224.251.0.0	255.255.224.0	proxy-arp 66																																																						
225.69.118.64	255.255.255.224	proxy-arp 66																																																						
225.97.128.0	255.255.128.0	proxy-arp 66																																																						
225.128.0.0	255.128.0.0	proxy-arp 66																																																						
225.134.0.0	255.255.192.0	proxy-arp 66																																																						
225.234.0.0	255.255.0.0	proxy-arp 66																																																						
225.234.207.128	255.255.255.128	proxy-arp 66																																																						
226.0.0.0	255.128.0.0	proxy-arp 66																																																						
226.14.9.20	255.255.255.254	proxy-arp 66																																																						
226.32.0.0	255.254.0.0	proxy-arp 66																																																						
226.108.76.0	255.255.254.0	proxy-arp 66																																																						
...																																																								
Related Commands																																																								
Note	This command is only available from the VIP address																																																							

Appendix A: MEX6200 System

A.1 MEX6200 Overview

The MEX6200 is a 3R WDM transponder, available as part of the MetroX® MTX62x0 family. It can multiplex several client signals (services) on a single fiber, each with a different wavelength, and transport them over a long distance. It is typically deployed as customer premises equipment (CPE) in enterprise campus environments and in central offices.

MEX6200 is available with 16 link ports forming up to 8 high-speed transponders that can serve lower-rate client signals of 10Gb/s.

MEX6200 is a highly integrated device that can incorporate MUX/DEMUX, Erbium Doped Fiber Amplifier (EDFA), and Dispersion Compensation Module (DCM) modules, to ensure data transmission to distances of up to 80km.

MEX6200 provides uplink 1+1 facility protection for the line ports. It also supports an optional Optical Switch that provides 1+1 facility protection for point-to-point topology.


A.2 Getting Started

This section describes the prerequisites and the procedure required to get the MetroX system.



Be sure to follow the installation instructions in the Quick Start Guide packaged within the MetroX bundle's box.

Figure 51: MetroX Connectivity to Switch



Be sure not to close TCP ports 81 and 82 in your network. Those ports are used for HTTP and HTTPs connection.

➤ **To access the MEX6200 webpage:**

Step 1. Navigate to the MetroX's WebUI.

Step 2. Go to Ports tab.

Step 3. Click the MEX6200 sidebar tab to reach the MEX6200 WebUI page.


Figure 52: MetroX Ports Tab Sidebar



Step 4. Enable the needed link ports via the WebUI by going to each port's configuration tab and clicking on the "Admin Up" button.

The following figure shows an example of the buttons used for performing operations in the Web application.

Figure 53: MEX6200 Item Buttons




The buttons displayed vary according to the configuration. For example, if MEX6200 does not have an EDFA module installed, the EDFA button is disabled.

The Item buttons displayed also vary according to the context of the window. For example, the FAN button is disabled in the Fault window because no faults are defined for this unit.

On the left side of the WebUI you can find the sidebar buttons.

Figure 54: Sidebar Buttons



A.3 Fault Management

This chapter describes the fault management of MEX6200, which is used to localize and identify problems in the network incorporating WDM extension units.

A.3.1 Alarms

MEX6200 keeps a list of the alarms currently detected on the system. When an alarm is detected, an “Alarm Rise” event is generated and the alarm is added to the list. When an “Alarm Clear” is detected, the alarm is removed from the list.

The following information is stored for each alarm:

- Date and Time – the date and time when the alarm was detected.
- Source – entity that caused the alarm.
- Severity – severity of the alarm.
- Type – type of the alarm.
- Service Affecting – Yes or No according to the alarm impact.

A.3.2 Events

MEX6200 continuously monitors the traffic signals and other exceptional conditions. Whenever such conditions occur, MEX6200 generates a time stamped event message and sends it as an SNMP notification to the registered management systems. MEX6200 logs the history of the last 512 events in a cyclic buffer that can be browsed by the Web application or by SNMP management systems.

In addition, the events and audit messages are printed in the system log files of MEX6200, which can be exported to a text file for offline viewing.

MEX6200 provides the following events:

- Alarm Rise: Alarms are standing faults. They are raised after a configurable stabilization period of several seconds. These events are generated when a new alarm occurs.
- Alarm Clear: Alarms are standing faults. They are cleared after a configurable stabilization period of several seconds. These events are generated when an alarm is cleared.

- **Link Up:** These are standard SNMP events that are generated when the operational status of a port is changed from Down to Up.
- **Link Down:** These are standard SNMP events that are generated when the operational status of a port is changed from Up to Down.
- **Cold Restart:** These are standard SNMP events that are generated after a Cold Restart to the node.
- **Warm Restart:** These are standard SNMP events that are generated after a Warm Restart to the node.
- **Test Status Changed:** These events are generated when the loopback or PRBS test status of a port is changed.
- **Protection Switching Event:** These events are generated when protection switching occurs.
- **Inventory Change:** These events are generated when the node inventory is changed.
- **Unsolicited Event:** These events are generated when an exceptional event occurs.
- **Configuration Change:** These events are generated when the node configuration is changed.

The MEX6200 generates an event when the configuration of a node is explicitly changed by the user and stores the event in the Configuration Changes log for auditing.

A.3.3 Alarm and Event Messages

This appendix describes the possible alarm and event messages.

The following table lists the MEX6200's alarm messages and explains their interpretation and/or corrective measures.

Table 69 - Alarm Messages

Source	Message	Interpretation/Corrective Measures
PSU1/ PSU2	Power Supply Failure	Replace the faulty PSU.
PSU1/ PSU2	Power Failure – Low Voltage	Replace the faulty PSU.
FAN	Fan Failure	The internal cooling fan of the device does not operate. Replace the FAN unit as soon as possible.
System	Hardware Failure	A technical failure has been detected. Replace the device.
System	Database Restore Failed	Failed to update the system configuration.
System	Database Restore in Progress	Failed to update the system configuration.
System	Cold Restart Required: FPGA Changed	After a warm restart, the FPGA version is not consistent with the software version. A cold restart is required.

Table 69 - Alarm Messages

Source	Message	Interpretation/Corrective Measures
System	Software Upgrade Failed	The downloaded software is corrupted. Reload the software.
System	Network Time Protocol Failure	SNTP timing protocol failure. Check the IP connection to the NTP servers.
External Input Alarm	(As configured)	The External Input Alarm is active.
GbE (Copper)	Ethernet Link Failure	Check the Auto Negotiation parameters.
Ethernet	Loss of Synchronization	Loss of Synchronization has been detected on the data links. Check that the input signal rate is correct.
Optics	Optics Removed	The optical module has been removed. Insert an optical module or shut the port down.
Optics	Optics Loss of Light	A Loss of Light indication has been received in regards to the specific optical module. The optical power of the received signal is below the minimum power level. Check the fiber connection and/or clean the fiber connector.
Optics	Optics Transmission Fault	The transceiver is not transmitting. Replace the optical module.
Optics	Optics Hardware Failure	A hardware fault was detected in the optical module. Replace the optical module.
Optics	Optics TX Loss of Lock	TX CDR Loss of lock.
Optics	Optics High Transmission Power	The transmission power of the optical module is above its specification.
Optics	Optics Low Transmission Power	The transmission power of the optical module is below its specification.
Optics	Optics High Temperature	The temperature inside the optical module is above its specification.
Optics	Optics Low Temperature	The temperature inside the optical module is below its specification.
Optics	Optics High Reception Power	The incoming signal into the optical module is too high. An attenuation of the input signal is required.
Optics	Optics Low Reception Power	The incoming signal into the optical module is too low.
Optics	Optics High Temperature	The temperature inside the optical module is above its specification.
Optics	Optics Low Temperature	The temperature inside the optical module is below its specification.

Table 69 - Alarm Messages

Source	Message	Interpretation/Corrective Measures
Optics	Optics High Laser Wave-length	The laser wavelength exceeds the high alarm level.
Optics	Optics Low Laser Wave-length	The laser wavelength exceeds the low alarm level.
Optics	Optics Loss Propagation	The laser was shut down due to a problem on the interface of the port mate.
Optics	Optics Bit Rate Mismatch	The inserted optical module has a mismatch problem due to the wrong rate or type. Replace the optical module or update the configured service type.
Optics	Unauthorized Optics Inserted and is Shutdown	The inserted optical module is unauthorized for use. Replace the optical module with an authorized optical module.
EDFA	EDFA Gain	The EDFA gain is out of acceptable range.
EDFA	EDFA Hardware failure	The interface does not respond.
EDFA	EDFA Temperature	The EDFA temperature is out of acceptable range.
EDFA	EDFA Loss of Light	No signal is detected.
EDFA	EDFA Receive Power Out of Bound	The receive signal is out of acceptable range. Check the optical power of the EDFA client signals. Use attenuation if required.
EDFA	EDFA Transmit Power Out of Bound	The transmit signal is out of acceptable range. Check the optical power of the EDFA client signals.
EDFA	EDFA Down	Closed the EDFA output upon loss of input. Check the EDFA client signals.
EDFA	EDFA Eye Safety	Hazard. No fiber is connected to the port.

The following table lists the configuration event messages generated by the MEX6200 and explains their interpretation.

Table 70 - Configuration Event Messages

Source	Message	Interpretation
System	Change date	The system date or time has changed.
System	Restore provisioning	A new configuration file has been loaded.
System	Change IP	The IP of the node has changed.
System	Configuration change	The system configuration was changed.
System	Alarm cut-off	The Alarm Cut-off has been operated.
System	Add user	A new user was added.
System	Delete user	A user was deleted.

Table 70 - Configuration Event Messages

Source	Message	Interpretation
System	Delete routing entry	A routing entry was deleted from the system Static Routing table.
System	Software Upgrade	Software Upgrade has been performed.
Port	Admin Down	Admin Down has been performed for the port.
Port	Admin Up	Admin Up has been performed for the port.
Link Port	Provisioning change	The provisioning of the port has changed.
Link Port	Test Operated	A test has been operated.
Link Port	Facility Loopback Released	A test has been released.
Link Port	Reset PM counters	Performance monitoring counters have been reset.
Service Port	Create APS	An APS was created for the service port.
Service Port	Remove APS	The APS for the service port has been removed.
Service Port	APS command	An APS command was issued.
Service Port	APS clear command	An APS command was cleared.

The following table lists the other event messages generated by MEX6200 and explains their interpretation.

Table 71 - Other Event Messages

Event Type	Source	Message	Interpretation
Inventory Changed	PSU, FAN, Optics	Inventory Changed	The node inventory has changed. A component was inserted or removed.
Switchover	Port	APS Switch Over	A protection switching event has occurred.
Test	Link Port	Test Mode changed	The link port test mode has changed.
ALS Status Changed	Port	ALS Laser occurred	The automatic laser shutdown was activated/deactivated.
Optical Power Drop	Link port	Power Level Drop	The Rx power of the port has been dropped by more than 2 dB since last interval.
Dying Gasp	System	Remote Unit Power Failure occurred	A remote unit had a power failure.

Table 71 - Other Event Messages

Event Type	Source	Message	Interpretation
Software Upgrade	System	Software Upgrade occurred	The software upgrade operation has been completed.

A.4 Configuration Management



When the two sides of MetroX are running different speeds, traffic bursts may occur which may lead to increased buffer usage on the side with the lower speed. Therefore, it is advised to add a switch in between the low speed cluster and the MetroX platform to absorb this traffic burst and align the speeds between two MetroX platforms.

Table 72 provides the possible means of configuring the MEX6200.

Table 72 - Configuration Options of the MEX6200


Tab	Description
System Configuration	Use the System Configuration window to configure general system parameters, view system inventory, configure SNTP parameters, configure IP addresses and static routing, configure SNMP parameters and traps and configure syslog servers.
Link Port Configuration	Use the Link Port Configuration window to configure the link port, configure the SFP/SFP+ module, configure ALS for a link port, and configure the APS for a link port (not available when the optional optical switch module is installed).
Management Port Configuration Gray	Use the Management Port Configuration window to configure management ports and enable/disable them, configure the SFP module, and configure ALS for a management port.
Ethernet Port Configuration	Use the Ethernet Port Configuration window to configure the Ethernet port status and parameters.
MUX/DEMUX Configuration Only see MUX	Use the MUX/DEMUX Configuration window to display the wavelengths of the uplink channels of MEX6200.
EDFA Configuration Gray	Use the EDFA Configuration window to configure the EDFA module and enable/disable the module.
COM Port Configuration Gray	Use the COM Port Configuration window to configure a COM port and enable/disable the port, configure APS for a COM port.
PSU Configuration	Use the PSU Configuration window to view information about the power supply units currently installed in the system.
Fan Unit Configuration	Use the Fan Unit Configuration window to view information about the fan unit currently installed in the system.

A.5 Performance Monitoring

MEX6200 provides port performance monitoring for link ports 1-16.

Use the Link Port Performance Monitoring window to view link port performance monitoring.

Figure 55: Link Port Performance Monitoring Window



➤ **To open the Link Port Performance Monitoring window:**

Step 1. Click Performance.

Step 2. Click a Port button to select the port.

The appropriate Link Port Performance Monitoring window opens.

A.5.1 Native Signal

Use the Link Port Performance Monitoring tab to view link port native signal performance monitoring.

Figure 56: Native Signal Performance Monitoring



➤ **To view native signal performance monitoring:**

Step 1. Click a Port button to select the link port.

The appropriate Link Port Performance Monitoring tab opens displaying the link port performance monitoring. The fields are explained in the following table. The counters are read only.

Step 2. From the PM Period drop-down list, select the interval.

Step 3. From the Type drop-down list, select Native Signal.

Step 4. Click Get PM.

The performance monitoring counters are updated.

Step 5. Export the PM information to a file:

1. Click Export to File.

The Opening table.csv dialog box appears.

2. Click Save File.

3. Click OK.

Step 6. To set the refresh rate of the PM display:

1. In the Refresh every field, type the number of seconds that the window should refresh.

The minimum refresh rate is 2 seconds.

1. Click Start Refresh.

The information is automatically updated after the specified number of seconds.

Step 7. Refresh the PM display manually. Click Refresh.

The information is updated immediately.

Step 8. Stop the automatic refresh of the PM display. Click Stop Refresh.

The automatic refresh is stopped and the Refresh every field is cleared.

Step 9. Clear the PM counters for a specific PM interval. In the table, at the end of the interval row, click Reset.

Step 10. Clear PM counters for a specific port. Click Reset Port PM.

Step 11. Clear PM counters for all ports. Click Reset All Ports PM.

Table 73 - Link Port Performance Monitoring Tab Parameters

Parameter	Description	Format/Values
PM Period	The interval for accumulating and displaying the performance monitoring counters.	15 Minutes, Days
Interval	The date and time of the interval.	PM Period is set to 15 Minutes: • Current: Performance monitoring counters accumulated during the current interval of 15 minutes are displayed in the first row. • 1 to 32: Performance monitoring counters accumulated during the last 32 intervals of 15 minutes are displayed in the second row to the last row of the table. PM Period is set to Days: • Untimed: Performance monitoring counters accumulated since last reset of the system or since the last reset of the performance monitoring counters are displayed in the first row of the table. • Current Day: Performance monitoring counters accumulated since 00:00 AM of the current day are displayed in the second row of the table. • Previous Day: Performance monitoring counters accumulated during the 24 hours since 00:00 AM of the previous day are displayed in the last row of the table.
Errors: • Coding Violation (CV) - or • B1 errors	The number of coding violation or B1 errors.	The number of 64B/66B coding violation errors detected during the performance monitoring interval. Note: This counter is service dependent.
Errored Seconds (ES)	The number of seconds in which at least one coding error was detected.	Number of seconds


Table 73 - Link Port Performance Monitoring Tab Parameters

Parameter	Description	Format/Values
Severely Errored Seconds (SES)	The number of seconds in which the number of errors detected crossed the threshold.	Number of seconds. Note: The counter stops when one of the following occurs: - The number of errors detected during the last second is below the threshold. - The Unavailable Seconds counter is incremented.
- Unavailable Seconds (UAS) or - Severely Errored Frames (SEF) or - Out of Frame seconds (OOF)	The number of unavailable seconds, severely errored frames, or out of frame seconds. ²	The count of Unavailable Seconds is incremented if the number of errors crossed the Severely Errored Seconds threshold at any time during the last 10 consecutive seconds. Note: This counter is service dependent.
Valid	Whether or not the performance monitoring interval has been completed, and whether or not the information is accurate.	- Partial: The measured interval has not been completed. - Yes: The performance monitoring interval has been completed. - No: The interval has been completed, but the performance monitoring information may not be accurate. Note: The performance monitoring information may be inaccurate due to one of the following reasons: - The performance monitoring counters of the interval were reset. - The node was reset during the interval. - The port was set to Admin Down during the interval. - The calendar time of the node was changed during the interval.

A.5.2 Optical Level

Use the Link Port Performance Monitoring tab to view link port optical level performance monitoring.

Figure 57: Optical Level Performance Monitoring



➤ **To view optical level performance monitoring:**

Step 1. Click a Port button.

The appropriate Link Port Performance Monitoring tab opens displaying the displaying the link port performance monitoring. The fields are explained in the following table. The counters are read only.

Step 2. From the PM Period drop-down list, select the interval.

Step 3. From the Type drop-down list, select Optical Level.

Step 4. Click Get PM.

The optical level counters are updated.

Step 5. To export the optical level information to a file:

1. Click Export to File.

The Opening table.csv dialog box appears.

2.Click Save File.

3.Click OK.

Step 6. To set the refresh rate of the PM display:

1. In the Refresh every field, type the number of seconds that the window should refresh.

The minimum refresh rate is 2 seconds.

2.Click Start Refresh.

The information is automatically updated after the specified number of seconds.

Step 7. To refresh the PM display manually, click Refresh .

The information is updated immediately.

Step 8. To stop the automatic refresh of the PM display, click Stop Refresh.

The automatic refresh is stopped and the Refresh every field is cleared.

Step 9. To clear the optical level counters for a specific port, click Reset Port PM.

Step 10. To clear the optical level counters for all ports, click Reset All Ports PM.

Table 74 - Link Port Performance Monitoring Tab Parameters

Parameter	Description	Format/Values
PM Period	The interval for averaging the measured Rx power.	15 Minutes, Days
Type	The type of performance monitoring.	Optical Level
Interval	The date and time of the interval.	PM Period is set to 15 Minutes: • Current: The date and time of the current interval of 15 minutes is displayed in the first row. • 1 to 32: The date and time of the last 32 intervals of 15 minutes is displayed in the second row to the last row of the table. PM Period is set to Days: • Untimed: The date and time of the last reset of the system or last reset of the optical level counters is displayed in the first row of the table. • Current Day: The date and 00:00 AM of the current day is displayed in the second row of the table. • Previous Day: The date and 00:00 AM of the previous day is displayed in the last row of the table.

Table 74 - Link Port Performance Monitoring Tab Parameters

Parameter	Description	Format/Values
Rx Level dBm	The measured Rx power level during the interval (in dBm).	PM Period is set to 15 Minutes: - Current: The measured Rx power for the current interval of 15 minutes is displayed in the first row. 1 to 32: The measured Rx power for the last 32 intervals of 15 minutes is displayed in the second row to the last row of the table. PM Period is set to Days: - Untimed: The average of the measured Rx power since last reset of the system or since the last reset of the optical level counters is displayed in the first row of the table. - Current Day: The average of the measured Rx power since 00:00 AM of the current day is displayed in the second row of the table. - Previous Day: The average of the measured Rx power during the 24 hours since 00:00 AM of the previous day is displayed in the last row of the table.

A.6 Maintenance

Use the maintenance tab to perform maintenance operations such as software upgrade, configuration management, restart and diagnostics.

Table 75 describes the available system maintenance tasks.

Table 75 - System Maintenance Options of the MEX6200

Tab	Description
Restart tab	Use the Restart tab to Cold Restart (for major upgrade to the device software), Warm Restart (for minor upgrade of the device software), or restore to factory settings.
Log Files tab	Use the Log Files tab to view and save System Log files.
Configuration tab	Use the Configuration tab to update the system configuration with a previously saved file of system configuration, while preserving or replacing the IP addresses, or to upload the current system configuration of MEX6200 and save it to the local file system.
Software tab	Use the Software tab to download software and to switch and activate a new software version.

Use the Link Port Maintenance window to perform diagnostic tests on link ports.

Use the External Alarm tab to configure the external alarm.

A.6.1 Upgrading Software on the MEX6200

Software upgrade on the MEX6200 system is not performed automatically with regular MLNX-OS® upgrade, and must be manually run as follows.

➤ **To perform software upgrade:**

- Step 1.** Click the “System” button at the top left of the WebUI page of MEX6200.
- Step 2.** Click the “Maintenance” sidebar button.
- Step 3.** Click the “Software” tab.
- Step 4.** Click the “Browse” button and locate the image to upgrade the software.
- Step 5.** Click “Download”.
- Step 6.** Reboot the MEX6200 by clicking “Switch and Cold Restart”.

Figure 58: MEX6200 Software Upgrade Webpage



Appendix B: Enhancing System Security According to NIST SP 800-131A

B.1 Overview

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.

B.2 Web Certificate

Mellanox supports signature generation of sha256WithRSAEncryption, sha1WithRSAEncryption self-signed certificates, and importing certificates as text in PEM format.

➤ **To configure a default certificate:**

Step 1. Create a new sha256 certificate. Run:

```
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` in the MLNX-OS User Manual.

Step 2. Show crypto certificate detail. Run:

```
switch (config) # show crypto certificate detail
```

Search for “signature algorithm” in the output.

Step 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:**

Step 1. Define the default hash algorithm. Run:

```
switch (config) # crypto certificate generation default hash-algorithm sha256
```

Step 2. Generate a new certificate with default values. Run:

```
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. Please refer to the MLNX-OS User Manual for further details.



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”.

B.3 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.
For more information please refer to the MLNX-OS User Manual.



SNMPv1 and SNMPv2 are not considered to be secure. To run in strict mode, only use SNMPv3.

B.4 SSH

The SSH server on the switch by default uses secure and unsecure ciphers, 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:**

```
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
- arcfour
- blowfish-cbc
- cast128-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.

B.5 HTTPS

By default, Mellanox switch supports HTTPS encryption using TLS1.0 up to TLS1.2. To work in strict mode you must configure the system to use TLS1.2. 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
switch (config)#
```

On top of enabling HTTPS, to prevent security breaches HTTP must be disabled.

- **To disable HTTP, run:**

```
switch (config)# no web http enable
```

B.6 LDAP

By default, Mellanox switch supports 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:**

```
switch (config) # ldap ssl mode {start-tls | ssl}
```



Both modes operate using SSL. The difference lies in the connection initialization and the port used.

Appendix C: Show Commands Supported by JSON API

Table 76 lists the “show” commands which currently support JSON API.



The character “*” indicates a wildcard.

Table 76 - JSON API Show Commands

Ethernet Commands
show interfaces ethernet transceiver
show interfaces ethernet * transceiver
show interface ethernet * status
show interface ethernet status
show interfaces ethernet * counters
show interfaces port-channel * counters
show interfaces mlag-port-channel * counters
show interfaces ethernet capabilities
show ip interface brief
show ip load-sharing
show lacp interfaces port-channel * system-identifier
show lldp timers
show lldp statistics global
show qos interface port-channel *
show vrf
show vrf *
show routing-context vrf
IB Commands
show interfaces ib transceiver
show interfaces ib * transceiver
show interface ib * status
show interface ib status
VPI Commands
show ports type
Chassis Management Commands
show leds

Table 76 - JSON API Show Commands

show leds *
show battery-backup-unit * detail
show asic-version
show bios
show inventory
show fan
show module
show system capabilities
show system profile
show version
General Commands
show ha dns
show files system
show stats sample
show stats sample *
show stats chd
show stats chd *
show stats cpu
show whoami
show users
show usernames
show xml-gw
show services small-servers
show ftp-server
show telnet-server