



Mellanox MLNX-OS® User Manual for Lenovo 90Y3474

Rev 4.20

Software Version 3.4.3002

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 NONINFRINGEMENT 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

Mellanox Technologies, Ltd.
Hakidma 26
Ofer Industrial Park
Yokneam 2069200
Israel
www.mellanox.com
Tel: +972 (0)74 723 7200
Fax: +972 (0)4 959 3245

© Copyright 2015. Mellanox Technologies. All Rights Reserved.

Mellanox®, Mellanox logo, BridgeX®, ConnectX®, Connect-IB®, CoolBox®, CORE-Direct®, GPUDirect®, InfiniBridge®, InfiniHost®, InfiniScale®, Kotura®, Kotura logo, Mellanox Connect. Accelerate. Outperform logo, Mellanox Federal Systems®, Mellanox Open Ethernet®, Mellanox Virtual Modular Switch®, MetroX®, MetroDX®, MLNX-OS®, Open Ethernet logo, PhyX®, ScalableHPC®, SwitchX®, TestX®, The Generation of Open Ethernet logo, UFM®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies, Ltd.

CyPU™, ExtendX™, FabricIT™, FPGADirect™, HPC-X™, Mellanox Care™, Mellanox CloudX™, Mellanox NEO™, Mellanox Open Ethernet™, Mellanox PeerDirect™, NVMeDirect™, StPU™, Switch-IB™, Unbreakable-Link™ are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

Table of Contents

Document Revision History	12
About this Manual	21
Chapter 1 Introduction	24
1.1 System Features	24
1.2 Ethernet Features	25
1.3 Gateway Features	25
Chapter 2 Getting Started	27
2.1 Configuring the Switch for the First Time	27
2.2 Starting the Command Line (CLI)	28
2.3 Starting the Web User Interface	28
2.4 Licenses	31
2.4.1 Installing MLNX-OS® License (CLI)	31
2.4.2 Installing MLNX-OS License (Web)	32
2.4.3 Retrieving a Lost License Key	34
2.4.4 Commands	36
Chapter 3 User Interfaces	41
3.1 Command Line Interface (CLI) Overview	41
3.1.1 CLI Modes	41
3.1.2 Syntax Conventions	42
3.1.3 Getting Help	42
3.1.4 Prompt and Response Conventions	43
3.1.5 Using the “no” Form	44
3.1.6 Parameter Key	45
3.1.7 Command Output Filtering	46
3.2 Web Interface Overview	47
3.2.1 Setup Menu	48
3.2.2 System Menu	49
3.2.3 Security Menu	50
3.2.4 Ports Menu	50
3.2.5 Status Menu	51
3.2.6 ETH Mgmt	51
3.2.7 IP Route	52
3.3 Secure Shell (SSH)	52
3.3.1 Adding a Host and Providing an SSH Key	52
3.3.2 Retrieving Return Codes when Executing Remote Commands	53
3.4 Commands	54
3.4.1 CLI Session	54
3.4.2 Banner	62
3.4.3 SSH	67
3.4.4 Remote Login	82
3.4.5 Web Interface	85

Chapter 4	System Management	100
4.1	Management Interface	100
4.1.1	Configuring Management Interfaces with Static IP Addresses	100
4.1.2	Configuring IPv6 Address on the Management Interface	100
4.1.3	Dynamic Host Configuration Protocol (DHCP)	100
4.1.4	Default Gateway	101
4.1.5	In-Band Management	101
4.1.6	Commands	103
4.2	NTP, Clock & Time Zones	149
4.2.1	Commands	150
4.3	Software Management	156
4.3.1	Upgrading MLNX-OS Software – Preconditions	156
4.3.2	Upgrading MLNX-OS Software	156
4.3.3	Upgrading MLNX-OS HA Groups	160
4.3.4	Deleting Unused Images	160
4.3.5	Downgrading MLNX-OS Software	161
4.3.6	Upgrading System Firmware	164
4.3.7	Image Maintenance via Mellanox ONIE	165
4.3.8	Commands	167
4.4	Configuration Management	178
4.4.1	Saving a Configuration File	178
4.4.2	Loading a Configuration File	178
4.4.3	Managing Configuration Files	178
4.4.4	Commands	181
4.5	Logging	207
4.5.1	Monitor	207
4.5.2	Remote Logging	207
4.5.3	Switch Power-On Self-Test	207
4.5.4	Commands	209
4.6	Debugging	227
4.6.1	Commands	228
4.7	Event Notifications	247
4.7.1	Supported Events	247
4.7.2	Terminal Notifications	248
4.7.3	Email Notifications	249
4.7.4	Commands	251
4.8	mDNS	270
4.8.1	Commands	271
4.9	User Management and Security	272
4.9.1	User Accounts	272
4.9.2	Authentication, Authorization and Accounting (AAA)	272
4.9.3	Commands	274
4.10	Cryptographic (X.509, IPSec)	315
4.10.1	Commands	316
4.11	Scheduled Jobs	328

4.11.1	Commands	328
4.12	Statistics and Alarms	338
4.12.1	Commands	338
4.12.2	Power Management	355
4.12.3	System Reboot	356
4.12.4	Commands	358
4.13	Network Management Interfaces	382
4.13.1	XML API	382
4.13.2	Commands	383
Chapter 5	Ethernet Switching	400
5.1	Interface	400
5.1.1	Break-Out Cables	400
5.1.2	Transceiver Information	403
5.1.3	Commands	404
5.2	Link Aggregation Group (LAG)	419
5.2.1	Configuring Static Link Aggregation Group (LAG)	419
5.2.2	Configuring Link Aggregation Control Protocol (LACP)	419
5.2.3	Commands	421
5.3	MLAG	439
5.3.1	MLAG Keepalive and Failover	441
5.3.2	Unicast and Multicast Sync	441
5.3.3	MLAG Port Sync	441
5.3.4	MLAG Virtual System-MAC	441
5.3.5	Upgrading MLAG Pair	441
5.3.6	MLAG Configuration	442
5.3.7	Commands	446
5.4	VLANs	463
5.4.1	Configuring Access Mode and Assigning Port VLAN ID (PVID)	463
5.4.2	Configuring Hybrid Mode and Assigning Port VLAN ID (PVID)	464
5.4.3	Configuring Trunk Mode VLAN Membership	464
5.4.4	Configuring Hybrid Mode VLAN Membership	465
5.4.5	Commands	466
5.5	QinQ	475
5.5.1	QinQ Operation Modes	475
5.5.2	Configuring QinQ	475
5.5.3	Commands	478
5.6	MAC Address Table	479
5.6.1	Configuring Unicast Static MAC Address	479
5.6.2	MAC Learning Considerations	479
5.6.3	Commands	480
5.7	Spanning Tree	486
5.7.1	Port Priority and Cost	486
5.7.2	Port Type	486
5.7.3	BPDU Filter	487
5.7.4	BPDU Guard	487

5.7.5	Loop Guard	487
5.7.6	Root Guard	488
5.7.7	MSTP	488
5.7.8	RPVST	488
5.7.9	Commands	491
5.8	OpenFlow	520
5.8.1	Flow Table	520
5.8.2	Configuring OpenFlow	521
5.8.3	Commands	522
5.9	IGMP Snooping	529
5.9.1	Configuring IGMP Snooping	529
5.9.2	Defining a Multicast Router Port on a VLAN	529
5.9.3	IGMP Snooping Querier	531
5.9.4	Commands	532
5.10	Link Layer Discovery Protocol (LLDP)	548
5.10.1	Configuring LLDP	548
5.10.2	DCBX	548
5.10.3	Commands	550
5.11	Quality of Service (QoS)	566
5.11.1	Priority Flow Control and Link Level Flow Control	566
5.11.2	Enhanced Transmission Selection (ETS)	567
5.11.3	Commands	570
5.12	Access Control List	580
5.12.1	Configuring Access Control List	580
5.12.2	ACL Actions	580
5.12.3	Commands	582
5.13	Port Mirroring	594
5.13.1	Mirroring Sessions	594
5.13.2	Configuring Mirroring Sessions	597
5.13.3	Verifying Mirroring Sessions	599
5.13.4	Commands	600
5.14	sFlow	609
5.14.1	Flow Samples	609
5.14.2	Statistical Samples	609
5.14.3	sFlow Datagrams	610
5.14.4	Sampled Interfaces	610
5.14.5	Configuring sFlow	610
5.14.6	Verifying sFlow	611
5.14.7	Commands	612
5.15	Transport Applications	624
5.15.1	RDMA over Converged Ethernet (RoCE)	624
5.16	802.1x Protocol	628
5.16.1	802.1x Operating Modes	628
5.16.2	Configuring 802.1x	629
5.16.3	Commands	630

Chapter 6	IP Routing	646
6.1	General	646
6.1.1	IP Interfaces	646
6.1.2	Equal Cost Multi-Path Routing (ECMP)	649
6.1.3	Virtual Routing and Forwarding	650
6.1.4	IPv4 Routing Mode	651
6.1.5	Commands	652
6.2	IPv6	692
6.2.1	Neighbor Discovery Protocol	692
6.2.2	Configuring IPv6	693
6.2.3	Commands	696
6.3	OSPF	722
6.3.1	Router ID	722
6.3.2	ECMP	722
6.3.3	Configuring OSPF	723
6.3.4	Verifying OSPF	724
6.3.5	Commands	727
6.4	BGP	762
6.4.1	State Machine	762
6.4.2	Configuring BGP	762
6.4.3	Verifying BGP	764
6.4.4	Commands	765
6.4.5	IP AS-Path Access-List	817
6.4.6	IP Community-List	819
6.5	Policy Rules	822
6.5.1	Route Map	822
6.5.2	IP Prefix-List	852
6.6	Multicast (IGMP and PIM)	855
6.6.1	Basic PIM-SM	855
6.6.2	Bidirectional PIM	856
6.6.3	PIM Load-Sharing	856
6.6.4	Bootstrap Router	857
6.6.5	Configuring Multicast	857
6.6.6	Commands	860
6.7	VRRP	902
6.7.1	Load Balancing	902
6.7.2	Configuring VRRP	903
6.7.3	Verifying VRRP	904
6.7.4	Commands	906
6.8	MAGP	917
6.8.1	MAGP Configuration	917
6.8.2	Commands	919
6.9	DHCP Relay	925
6.9.1	Commands	926
Appendix A	Enhancing System Security According to NIST SP 800-131A	934

	A.1 Overview	934
	A.2 Web Certificate.....	934
	A.3 Code Signing	935
	A.4 SNMP.....	935
	A.5 SSH.....	936
	A.6 HTTPS	936
	A.7 LDAP	937
	A.8 Password Hashing	939
Appendix B	Security Vulnerabilities and Exposures	940
Appendix C	UI Changes in Version 3.4.2008	956
	C.1 Interface Addressing Change	956
	C.2 CLI Change.....	956
	C.3 MIB ifTable Change	957
	C.4 WebUI Ports Page Change	958
	C.5 Interface Speed Configuration Change	958
	C.6 CLI Change.....	958
	C.7 WebUI Change	959
	C.8 IB SM Link Speed Change.....	960
	C.9 Multi-ASIC Support.....	961
	C.10 CLI Change.....	961
	C.11 MIB entPhysicalTable Change	963
	C.12 MGMT Module Display Change	964
	C.13 MLNX-OS Image Name Change.....	965
	C.14 CLI Change.....	965
	C.15 WebUI Status Page Change	965
	C.16 CPU Module Display Change	966
	C.17 CLI Change.....	966
	C.18 WebUI System Inventory Page Change	967

List of Tables

Table 1 -Reference Documents	21
Table 2 -Glossary	21
Table 3 -General System Features	24
Table 4 -Ethernet Features	25
Table 5 -Gateway Features	25
Table 6 -Serial Terminal Program Configuration for PPC Based Systems	27
Table 7 -CLI Modes and Config Context	41
Table 8 -Syntax Conventions	42
Table 9 -Angled Brackets Parameter Description	45
Table 10 -WebUI Setup Submenus	48
Table 11 -WebUI System Submenus	49
Table 12 -WebUI Security Submenus	50
Table 13 -WebUI Ports Submenus	50
Table 14 -WebUI Status Submenus	51
Table 15 -WebUI ETH Mgmt Submenus	51
Table 16 -WebUI IP Route Submenus	52
Table 17 -POST Return Codes	207
Table 18 -Supported Event Notifications and MIB Mapping	247
Table 19 -User Roles (Accounts) and Default Passwords	272
Table 20 -LWR Configuration Behavior	356
Table 21 -Key for Port Splitting Figure	401
Table 22 -Port Splitting Options	401
Table 23 -Supported VLANs by RPVST per Switch System	489
Table 24 -Mirroring Parameters	595
Table 25 -List of Statistical Counters	609
Table 26 -Common Vulnerabilities and Exposures	940

List of Figures

Figure 1:	Managing an Ethernet Fabric Using MLNX-OS	26
Figure 2:	IBM System Console Ports	27
Figure 3:	MLNX-OS Login Window	29
Figure 4:	EULA Prompt	30
Figure 5:	Welcome Popup	30
Figure 6:	Display After Login	31
Figure 7:	No Licenses Installed	32
Figure 8:	Enter License Key(s) in Text Box	33
Figure 9:	Installed License	34
Figure 10:	WebUI	48
Figure 11:	SX65xx Downgrade Attention Sticker	161
Figure 12:	Break-Out Cable	400
Figure 13:	Port Splitting Options	401
Figure 14:	Basic MLAG Setup	439
Figure 15:	Basic MLAG Topology	442
Figure 16:	MAC Learning Disable Example Case	479
Figure 17:	RPVST Network Config	488
Figure 18:	RPVST and RSTP Cluster	489
Figure 19:	Overview of Mirroring Functionality	594
Figure 20:	Mirror to Analyzer Mapping	594
Figure 21:	Header Format Options	597
Figure 22:	Mirroring Session	598
Figure 23:	sFlow Functionality Overview	609
Figure 24:	RoCEv2 and RoCE Frame Format Differences	624
Figure 25:	RoCEv2 Protocol Stack	625
Figure 26:	ECMP	649
Figure 27:	Multiple Hash Functions	650
Figure 28:	IPv6 Network	693
Figure 29:	OSPF Basic Topology	723
Figure 30:	Basic BGP Configuration	762
Figure 31:	Common VRRP Configuration with Load Balancing	902
Figure 32:	1U MIB ifTable Before Screenshot	957
Figure 33:	1U MIB ifTable After Screenshot	957
Figure 34:	Director Switch MIB ifTable Before Screenshot	958
Figure 35:	Director Switch MIB ifTable After Screenshot	958

Figure 36: Ports WebUI Page	960
Figure 37: MIB entPhysicalTable Before Screenshot	963
Figure 38: MIB entPhysicalTable After Screenshot	964
Figure 39: Status WebUI Page	966
Figure 40: System Inventory WebUI Page	967

Document Revision History

Rev 4.20 – July 30, 2015

Added:

- the command “switchport dot1q-tunnel qos-mode” on page 471
- Section 5.5, “QinQ,” on page 475
- the command “dot1x host-mode” on page 633
- the command “show ip route” on page 680
- the command “vlan-pop” on page 589
- the command “vlan-push” on page 590

Updated:

- Section 2.4, “Licenses,” on page 31
- the command “ssh server host-key” on page 68
- notes of the command “aaa authorization” on page 288
- the command “show module” on page 369
- the command “snmp-server user” on page 393
- the command “ip ospf authentication-key” on page 750
- the command “neighbor password” on page 792
- the command “neighbor peer-group” on page 793
- Appendix B, “Security Vulnerabilities and Exposures” on page 940

Rev 4.10 – June 11, 2015

Added:

- Section 2.1, “Configuring the Switch for the First Time,” on page 27 with MLNX-OS® Boot Menu step
- the command “ssh server security strict” on page 73
- the command “ssh server tcp-forwarding enable” on page 74
- Section 4.1.5, “In-Band Management,” on page 101
This feature can now be enabled with IP Routing. Also updated the flow of setting an in-band management channel.
- the command “show module” on page 369
- Section 5.1.1, “Break-Out Cables,” on page 400
- the command “ip address dhcp” on page 411
- the command “ip address dhcp” on page 428
- Section 5.3.4, “MLAG Virtual System-MAC,” on page 441
- Section 5.3.5, “Upgrading MLAG Pair,” on page 441
- Section 5.16, “802.1x Protocol,” on page 628
- Section 6.1.3, “Virtual Routing and Forwarding,” on page 650

- the command “ip l3” on page 652
- the command “vrf definition” on page 653
- the command “routing-context vrf” on page 654
- the command “description” on page 656
- the command “rd” on page 657
- the command “vrf forwarding” on page 658
- the command “show routing-context vrf” on page 660
- the command “show vrf” on page 661
- the command “ip address dhcp” on page 666
- Section 6.2, “IPv6,” on page 692 commands by adding loopback interface configuration mode to the commands
- Section 6.6.3, “PIM Load-Sharing,” on page 856
- the command “ip pim multipath rp” on page 874
- Appendix C, “UI Changes in Version 3.4.2008,” on page 956

Updated:

- the command “tcpdump” on page 148
- Section 4.3.2, “Upgrading MLNX-OS Software,” on page 156 with HA group note
- Section 4.3.3, “Upgrading MLNX-OS HA Groups,” on page 160
- the command “show inventory” on page 368
- the command “show asic-version” on page 371
- Section 5.3.1, “MLAG Keepalive and Failover,” on page 441
- Step 10 in Section 5.3.6, “MLAG Configuration,” on page 442
- the example of the command “upgrade-timeout” on page 457
- the command “ip routing” on page 655
- the command “show ip routing” on page 659
- the command “show ip interface” on page 673
- the command “interface loopback” on page 674 “id” parameter range
- the command “ip route” on page 678
- the command “show ip route” on page 680
- the command “clear ip arp” on page 685
- the command “show ip arp” on page 686
- the command “ping” on page 687
- the command “traceroute” on page 688
- the command “tcpdump” on page 690

Removed:

- the command “interface vlan create” from Section 4.1.6, “Commands,” on page 103
- the command “ipv6 dhcp client”

- Section B.9, “Security Vulnerabilities and Exposures,” on page 1330 and added it to the RN

Rev 3.70 – March 19, 2015

Updated:

- the command “[speed](#)” on page 1065
- the command “[show interfaces ib](#)” on page 1071
- the command “[show interfaces ib status](#)” on page 1072

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 456
- the command “[upgrade-timeout](#)” on page 457
- [Section 5.7.4, “BPDU Guard,”](#) on page 487

Updated:

- MLAG configuration verification [Step 1](#) with system MAC and upgrade timeout
- the command “[show mlag](#)” on page 458
- [Table 23, “Supported VLANs by RPVST per Switch System,”](#) on page 489

Rev 3.60 – March 05, 2015

No changes

Rev 3.50 – February 24, 2015

Added:

- the command “[show version concise](#)” on page 365

Updated:

- the command “[show uboot](#)” on page 366

Rev 3.40 – February 11, 2015

Added:

- “List of Tables” and “List of Figures” Sections
- Updated [Section 2.4, “Licenses,”](#) on page 31
- the command “[license delete](#)” on page 38
- the command “[license install](#)” on page 39
- the command “[telnet](#)” on page 82

- the command “terminal” on page 59
- the command “web cache-enable” on page 86
- the command “ip default-gateway” on page 105
- the command “boot system” on page 169
- the command “configuration write” on page 203
- the command “logging trap” on page 224
- the command “email autosupport enable” on page 251
- the command “email autosupport event” on page 252
- the command “crypto ipsec ike” on page 316
- the command “lacp-individual enable” on page 427
- the command “show interfaces port-channel” on page 435
- the command “show interfaces port-channel compatibility-parameters” on page 436
- the command “show interfaces port-channel load-balance” on page 437
- the command “show interfaces port-channel summary” on page 438
- Section 5.7.8, “RPVST,” on page 488
- the command “spanning-tree vlan forward-time” on page 510
- the command “spanning-tree vlan hello-time” on page 511
- the command “spanning-tree vlan max-age” on page 512
- the command “spanning-tree vlan priority” on page 513
- the command “show spanning-tree vlan” on page 519
- Section 6.2, “IPv6,” on page 692
- the command “auto-cost reference-bandwidth” on page 731
-
- the command “show ip multicast interface proxy-arp” on page 1326

Updated:

- Section 2.3, “Starting the Web User Interface,” on page 28
- the command “image options” on page 175
- the command “reload” on page 185
- Section 4.5.2, “Remote Logging,” on page 207
- the command “logging debug-files” on page 211
- Section 4.6.1, “Commands,” on page 228
- Section 4.9.1, “User Accounts,” on page 272
- the command “username” on page 274
- the command “aaa authentication attempts track” on page 281
- the command “radius-server host” on page 293
- the command “tacacs-server host” on page 297
- the command “snmp-server auto-refresh” on page 383
- the command “snmp-server user” on page 393

- the command “show interfaces ethernet [<inf>] description” on page 415
- the command “show interfaces ethernet [<inf>] status” on page 416
- the command “show interfaces port-channel summary” on page 438
- the command “show interfaces mlag-port-channel summary” on page 461
- the command “spanning-tree mode” on page 492
- the command “show spanning-tree” on page 514
- the command “show spanning-tree detail” on page 515
- the command “show spanning-tree interface” on page 516
- the command “show spanning-tree mst” on page 517
- the command “show spanning-tree root” on page 518
- Section 5.9.2, “Defining a Multicast Router Port on a VLAN,” on page 529
- the command “dcb application-priority” on page 557
- the command “dcb priority-flow-control enable” on page 576
- Section 5.14.1, “Flow Samples,” on page 609
- the command “ip arp timeout” on page 684
- the command “redistribute” on page 733

Rev 3.30 – November 19, 2014

Updated:

- the command “web https” on page 93
-
- the command “dcb application-priority” on page 557
- Section A.6, “HTTPS,” on page 936
- Section A.8, “Password Hashing,” on page 939

Rev 3.20 – November 09, 2014

Added:

- Section 5.6.2, “MAC Learning Considerations,” on page 479
- the command “mac-learning disable” on page 482
- Appendix A, “Enhancing System Security According to NIST SP 800-131A,” on page 934

Updated:

- Section 1.2, “Ethernet Features,” on page 25
- Section 3.2, “Web Interface Overview,” on page 47
- the command “interface port-channel” on page 421
- the command “show lacp interfaces neighbor” on page 431
- the command “ip icmp redirect” on page 672
- Section 6.4, “BGP,” on page 762

Replaced:

- the command “show lacp interfaces port-channel” with the command “show lacp” on [page 433](#)
- the command “show lacp system-identifier” with the command “show lacp interfaces system-identifier” on [page 434](#)

Rev 3.10 – July 20, 2014

Added:

- [Section 5.15, “Transport Applications,”](#) on [page 624](#)
- [Section 6.1.1, “IP Interfaces,”](#) on [page 646](#)

Updated:

- [Section 4.13.1, “XML API,”](#) on [page 382](#)
- [MAC addresses note in Section 5.3, “MLAG,”](#) on [page 439](#)

Rev 3.00 – June 05, 2014

No changes

Rev 2.90 – 19 May, 2014

Updated:

- the command “show configuration” on [page 205](#)
- the command “show uboot” on [page 366](#)
- the command “show voltage” on [page 375](#)

Rev 2.80 – May 08, 2014

Added:

- [supported versions note in Section 5.9, “IGMP Snooping,”](#) on [page 529](#)

Rev 2.70 – April 30, 2014

Added:

- [Appendix A, “Enhancing System Security According to NIST SP 800-131A,”](#) on [page 934](#)
- [supported versions note in Section 5.9, “IGMP Snooping,”](#) on [page 529](#)

Updated:

- the command “show ssh server” on [page 81](#)
- the command “web auto-logout” on [page 85](#)
- the command “web https” on [page 93](#)
- the command “show web” on [page 99](#)
- the command “show usernames” on [page 276](#)
- the command “ldap base-dn” on [page 300](#)

- the command “[ldap ssl](#)” on page 310

Rev 2.60 – April 10, 2014

N/A

Rev 2.50 – April 2014

Updated:

- [Section 3.1.7, “Command Output Filtering,”](#) on page 46
- the command “[show protocols](#)” on page 381
- [Section 5.1.1, “Break-Out Cables,”](#) on page 400
- the command “[show mac-address-table](#)” on page 484
- the command “[deny/permit \(MAC ACL rule\)](#)” on page 584
- the command “[show mac/ipv4 access-lists](#)” on page 592

Rev 2.40 – February, 2014

Updated:

- [Section 4.3.6.2, “Importing Firmware and Changing the Default Firmware,”](#) on page 165 – updated Step 1
- the command “[show running-config](#)” on page 206
- the command “[show log](#)” on page 226
- [Section 4.10, “Cryptographic \(X.509, IPSec\),”](#) on page 315
- [Section 5.2.1, “Configuring Static Link Aggregation Group \(LAG\),”](#) on page 419 – removed unnecessary step
- the command “[lldp tlv-select](#)” on page 556
- the command “[show lldp interface](#)” on page 559

Added:

- [Section 3.1.7, “Command Output Filtering,”](#) on page 46

Rev 2.30 – January, 2014

Updated:

- the command “[crypto certificate generation](#)” on page 321
- the command “[crypto certificate name](#)” on page 322

Rev 2.20 – January, 2014

N/A

Rev 2.10 – January, 2014

Added:

- [Section 4.12.2.1, “Width Reduction Power Saving,”](#) on page 355

Updated:

- Section 2.2, “Starting the Command Line (CLI),” on page 28
- Section 2.3, “Starting the Web User Interface,” on page 28
- Section 4.3.2, “Upgrading MLNX-OS Software,” on page 156 with EULA note
- the command “load-interval” on page 410 with Config Interface Port Channel
- the command “spanning-tree port-priority” on page 496 with Config Interface Port Channel
- Section 5.8, “OpenFlow,” on page 520
- the command “openflow description” on page 523
- the command “show openflow” on page 528
- the command “switchport {hybrid, trunk} allowed-vlan” on page 473 with Config Interface Port Channel
- the command “spanning-tree cost” on page 497 with Config Interface Port Channel
- the command “spanning-tree port type” on page 498 with Config Interface Port Channel
- the command “spanning-tree guard” on page 499 with Config Interface Port Channel
- the command “spanning-tree bpdupfilter” on page 500 with Config Interface Port Channel
- the command “deny/permit (IPv4 ACL rule)” on page 585
- the command “sflow enable (interface)” on page 622 with Config Interface Port Channel
- Section 6.3, “OSPF,” on page 722
- the command “router-id” on page 729

Rev 2.00 – December 2013

Added:

- Section 5.1.2, “Transceiver Information,” on page 403

Updated:

- Section 4.3.2, “Upgrading MLNX-OS Software,” on page 156
- Section 4.3.4, “Deleting Unused Images,” on page 160
- Section 4.6, “Debugging,” on page 227
- the example of the command “show cpld” on page 367
- “Notification Indicator” column in Section 8.3.2, “Standalone Proxy-ARP Configuration,” on page 1271
- the command “lldp tlv-select” on page 556

Moved:

Section 4.5.3, “Switch Power-On Self-Test,” on page 207 from 4.11.1

Section 3.3, “Secure Shell (SSH),” on page 52 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 1329

Rev 1.80 – October 2013

Added:

- [Section 5.7.7, “MSTP,”](#) on page 488
- [Section 5.8, “OpenFlow,”](#) on page 520
- [Section 5.9.3, “IGMP Snooping Querier,”](#) on page 531
- the command “ip igmp snooping querier”
- the command “igmp snooping querier query-interval”
- the command “show ip igmp snooping querier”
- [Section 5.10.2, “DCBX,”](#) on page 548
- the command “lldp tlv-select dcbx”
- the command “dcb application-priority”
- the command “show dcb application-priority”

Updated:

- the command “show lldp interface”
- the command “show lldp interfaces ethernet <inf> remote”

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.
Director switch Installation Guide	Each Mellanox Technologies' switch platform is shipped with an Installation Guide document to bring-up and initialize the switch platform.
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.

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.
ARP	Address Resolution Protocol. A protocol that translates IP addresses into MAC addresses for communication over a local area network (LAN).

Table 2 - Glossary

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.
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.
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
HA (High Availability)	A system design protocol that provides redundancy of system components, thus enables overcoming single or multiple failures in minimal downtime
Host	A computer platform executing an Operating System which may control one or more network adapters
LACP	Link Aggregation Control Protocol (LACP) provides a method to control the bundling of several physical ports together to form a single logical channel. LACP allows a network device to negotiate an automatic bundling of links by sending LACP packets to the peer (directly connected device that also implements LACP).
LDAP	The Lightweight Directory Access Protocol is an application protocol for reading and editing directories over an IP network.
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.

Table 2 - Glossary

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.
SNMP	Simple Network Management Protocol. A network protocol for the management of a network and the monitoring of network devices and their functions
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 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 controls
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

1.2 Ethernet Features


Table 4 - Ethernet Features

Feature	Description
General	• ACL – 24K rules (permit/deny) • Breakout cables • Jumbo Frames (9K)
Ethernet support	• 48K Unicast MAC addresses • 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 • 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 Gateway Features

Table 5 - Gateway Features

Feature	Description
Proxy-ARP	• Proxy-ARP interface • Unicast • Multicast • High availability Proxy-ARP

Figure 1: Managing an Ethernet Fabric Using MLNX-OS

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 (mini USB) port of the switch system using the supplied cable.

Figure 2: IBM System Console Ports



No remote IP connection is available at this stage via the external management port. The internal management port can be accessed currently by the chassis management.

Step 2. 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
Parity	None
Flow Control	None

Step 3. 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 4. Login as *admin* and use *admin* as password.

If the machine is still initializing, you might not be able to access the CLI until initialization completes. As an indication that initialization is ongoing, a countdown of the number of remaining modules to be configured is displayed in the following format: “<no. of modules> Modules are being configured”.

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

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

- 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: `https://<switch_IP_address>`.
- Step 4.** Login to the switch (default user name is *admin*, password *admin*).

Figure 3: MLNX-OS Login Window

Mellanox TECHNOLOGIES

Mellanox MLNX-OS Management Console

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

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

Login

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

Account

Password


Login

Mellanox MLNX-OS Switch Management . Best viewed using Firefox, Chrome, IE 7 or higher at 1024x768 resolution or higher.

© 2009-2012 Mellanox Technologies, Inc.

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

Figure 4: 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 5: 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 6.

Figure 6: 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.

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 7: No Licenses Installed

Mellanox MLNX-OS SX6506 Management Console
Host: switch-113dc8 User: admin Logout

Standalone Virtual IP Active node Chassis master Subnet Manager is not running.

Setup System Security Ports Status IB SM MGMT Fabric Inspectr ETH MGMT Save

Licensing

Interfaces
HA
Routing
DNS
Hostname
Hosts
ARP
Neighbors
Virtual Switch Mgmt
Web
SNMP
Email Alerts
XML gateway
Logs
Configurations
Date and Time
NTP
Licensing

System Serial Number
MXXXXXXXXXX7

Installed Licenses

License	Key	Feature	Valid	Max num ufm ports supported	Active
LK2-EFM_SX-5P26-85G2-3488-A3MG-VD3V-E7U	☐	EFM_SX	yes	200	yes

Remove

Add New License(s)

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

Add Licenses Save

© 2009-2012 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 8: Enter License Key(s) in Text Box

Mellanox MLNX-OS SX6506 Management Console
 Host: switch-113dc8 User: admin Logout
 Standalone Virtual IP Active node Chassis master Subnet Manager is not running.

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

Licensing

Interfaces
 HA
 Routing
 DNS
 Hostname
 Hosts
ARP
 Neighbors
 Virtual Switch Mgmt
 Web
 SNMP
 Email Alerts
 XML_gateway
 Logs
 Configurations
 Date and Time
 NTP
 Licensing

System Serial Number
 M*****7

Installed Licenses

License	
☐ Key	LK2-EFM_SX-5P26-85G2-3488-A3MG-VD3V-E7U
Feature	EFM_SX
Valid	yes
Max num ufm ports supported	200
Active	yes

Remove

Add New License(s)

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

<your license key>

Add Licenses Save

© 2009-2012 Mellanox Technologies, Inc.

All installed licenses should now be displayed.

Figure 9: Installed License

Mellanox MLNX-OS SX6506 Management Console
 Host: switch-113dc8 User: admin Logout
 Standalone Virtual IP Active node Chassis master Subnet Manager is not running.

Licensing

System Serial Number
 MXXXXXXXXXX7

Installed Licenses

License	Key	Feature	Valid	Max num ufm ports supported	Active
LK2-XXXXXXXXXX-E7U		EFM_SX	yes	200	yes

Add New License(s)

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

Add Licenses **Save**

© 2009-2012 Mellanox Technologies, Inc.

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                Type                Part number         Serial Number
=====
CHASSIS               SX1035              MSX6036F-1BFR       MT1121X02692
MGMT                  SX1035              MSX6036F-1BFR       MT1121X02692
FAN                   SXX0XX_FAN          MSX60-FF             MT1121X02722
PS1                   SXX0XX_PS           N/A                  N/A
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 31](#).

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_LEN_ETH_User_Manual.pdf <scp://username[:password]@hostname/path/filename> switch (config) #</pre>	
Related Commands		
Note		

license delete

license delete <license-key>

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	

license install

licenses 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	<pre>switch (config) # licenses install <license-key> switch (config) #</pre>
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 (CLI) Overview

MLNX-OS® is equipped with an industry-standard 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 7 - 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 loopback.
Config Interface Ethernet	Configuration mode for Ethernet interface.
Config Interface Port Channel	Configuration mode for Port channel (LAG).
Config VLAN	Configuration mode for VLAN.
Any Command Mode	Several commands such as “show” 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 8 - Syntax Conventions

Syntax Convention	Description	Example
< > Angled brackets	Indicate a value/variable that must be replaced.	<1...65535> or <switch inter-face>
[] 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 9 - 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.
<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.

Table 9 - Angled Brackets Parameter Description

Parameter	Description
<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 Command Output Filtering

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   Type           Present Power Is Fatal
=====
MGMT     SX1036             1         1    Not Fatal
FAN      SXX0XX_FAN         1         1    Not Fatal

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

```

3.2 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 10: WebUI

Mellanox MLNX-OS SX6036G Management Console
Host: switch-6287a4 User: admin Logout

Standalone Virtual IP Active node Subnet Manager is not running.

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

Ports Information Product Documents

Ports
Phy Profile
Monitor Session
Protocol Type

Port channels: 1 2

Port Info

Port number :	1	Mac address :	00:02:c9:72:0d:ac
Port type :	ETH	MTU :	1500 bytes
Port description :		Flow-control :	receive off send off
Admin state :	Enabled	Actual speed :	10 Gbps
Operational state :	Down	Switchport mode :	access
PFC admin mode :	Off	PFC operational mode :	Off
Last clearing of 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		

Port Counters Clear Port 1 Counters

RX packets :	0	TX packets :	0
RX unicast packets :	0	TX unicast packets :	0

3.2.1 Setup Menu

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

Table 10 - WebUI Setup Submenus

Submenu Title	Description
Interfaces	Obtains the status of, configures, or disables interfaces to the InfiniBand 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. Note: Changing hostname stamps a new HTTPS certificate.
DNS	Configures, removes, modifies or displays static and dynamic name servers.
Login Messages	Edits the login messages: Message of the Day (MOTD), Remote Login message, and Local Login message.

Table 10 - WebUI Setup Submenus

Submenu Title	Description
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 11 - 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.
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.

Table 11 - WebUI System Submenus

Submenu Title	Description
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 12 - 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 13 - 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.
Monitor Session	Displays monitor session summary and enables configuration of a selected session.

3.2.5 Status Menu

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

Table 14 - 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.

3.2.6 ETH Mgmt



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

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

Table 15 - WebUI ETH Mgmt Submenus

Submenu Title	Description
Spanning Tree	Configures and monitors spanning tree protocol.

Table 15 - WebUI ETH Mgmt Submenus

Submenu Title	Description
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.
Priority Flow Control	Manages priority flow control.

3.2.7 IP Route

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

Table 16 - 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 100 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
AAAAB3NzaC1yc2EAAAABIwAAAIEAsXeklqc8T0EN2mnMcVcfhueaRYzIVqt4rVsRERiJm1Jh4mkYYIa8hGGikNa+
t5xw2dRrNxnHYLK51bUsSG1ZNwZT1Dpme3pAZeMY7G4ZMgGIW9xOuaXgAA3eBeoUjFdi6+1BqchWk0nTb+gMfI/
MK/heQNns7AtTrvqg/O5ryIc="
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 Commands

3.4.1 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 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 localhost:0.0.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	

Example

```
switch (config) # cli session auto-logout  
switch (config) #
```

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	

3.4.2 Banner

banner login

banner {login | login-remote | login-local} <string>
no banner login

Sets the CLI welcome banner message. The login-remote refers to the SSH connections banner, while the login-local refers to the serial connection banner. 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	<pre>switch (config) # banner login example switch (config) # show banner Banners: MOTD: Mellanox Switch Login: example switch (config) #</pre>
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	
Role	admin	
Example	<pre>switch (config) # banner login-local Testing switch (config) #</pre>	
Related Commands	show banner	
Note	If more then 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	
Role	admin	
Example	<pre>switch (config) # banner login-remote Testing switch (config) #</pre>	
Related Commands	show banner	
Note	If more then 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

Sets 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	<pre>switch (config) # banner motd "My Banner" switch (config) # show banner Banners: MOTD: My-Banner Login: Mellanox MLNX-OS Switch Management switch (config) #</pre>	
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
Role	unpriv/monitor/admin
Example	<pre>switch (config) # show banner Banners: MOTD: Testing Login: Mellanox MLNX-OS Switch Management switch (config) #</pre>
Related Commands	<pre>banner login banner motd</pre>
Note	

3.4.3 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: 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	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-
IEArB9i5OnukAHNUOkwpCmEl0m88kJgBzL22+F5tfaSn+S0pVYxrcZeYuzXsoZ1VtFTk2-
FydwY0YvMS0Kcv2PuCrPZV/
GYd3lQEnn22rEmrlPrKCrMl1XlUy6DFlr3OgwWmlbaobmDlG/gSziWz/gc4Jgqf2CyX-
Fq4pzaRljarlVvk="
    DSA v2 host key: "switch-5ea5d8 ssh-dss AAAAB3NzaC1kc3MAAAC-
BAMeJ3S+nyaHhRbwv3tJqlWttDC35RZVC5iG4ZEvmMHp28VL94OcyuGh39VCdM9pEVaI7h
zZrsgHrNqakb/YLD/7anGH3wpl9Fx8lfe0RH3bloJzG+mJ6R5momdoPCrKwEKiKABKE00-
jLzlVznpP0IHxjwF+TbR3dK5HwVzQYw/bAAAAFQCBODPqBZZa+2KylK1zUsbZ2pKhgQAAA-
IAJK+StiQdtORw1B5UCMzTrTef5L07DSfVreMEYtTRnBBtgVSNqQFWpSQIYbVDHQR9T6qCM
4VO39DuHUGQ1TMDIX7t+9mfbB87YyUu5a/ndbf3GhNhXHWwbzlr9hgLL7FSHA7DYH7bVOZ-
RlqxH64eQKGZqylps/F4E31lyn7GC4EQAAAIa/2osHipXf+NRjplgfmHROVvf/mGE9Vzc9/
AMUxlJJn5VhvEJ5CZW9cI+LxMOJojhOj3YW3BlczGxRObDA9vUbKXTNc8bkgoUrxySAH1rH
N0PqJgeT4L009AItSp3mlmxHqds7jixfTvOTEKWxrgpczlmTB8+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 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	3.1.0000
Role	admin
Example	switch (config) # ssh server security strict switch (config) #
Related Commands	show ssh server
Note	

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	<pre>switch (config) # ssh client global host-key-check no switch (config) # ssh client global known-host "72.30.2.2 ssh-rsa AAAAB3NzaClyc2EAAAABIwAAAIEArB9i5OnukAHNUOkwpCmEl0m88kJgB- zL22+F5tfaSn+S0pVYxrceZeyuzXsoZlVtFTk2FydwY0YvMS0Kcv2PuCrPZV/ GYd3lQEnn22rEmrlPrKCrMl1XlUy6DFlr3OgwWm1baobmDlG/gSziWz/gc4Jgqf2CyX- Fq4pzaR1jarlVk=" 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

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@]host-name [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 estab- lished. 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
	3.3.5050 Updated Example
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 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	

3.4.4 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	<pre>switch (config) # telnet-server enable switch (config) # show telnet-server Telnet server enabled: yes</pre>
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	<pre>switch (config) # show telnet-server Telnet server enabled: yes switch (config) #</pre>
Related Commands	telnet-server enable
Note	

3.4.5 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 3.3.5050	Updated Example
Role	admin	
Example	<pre>switch (config) # web auto-logout 60 switch (config) # show web</pre> Web User Interface: <pre>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.</pre> <pre>Inactivity timeout: 1 hr Session timeout: 2 hr 30 min Session renewal: 30 min</pre> Web file transfer proxy: <pre>Proxy enabled: no</pre> Web file transfer certificate authority: <pre>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	<pre>switch (config) # web client cert-verify</pre>
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 3.3.5050 Updated Example
Role	admin
Example	<pre>switch (config) # web enable switch (config) # show web</pre> Web User Interface: <pre>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.</pre> <pre>Inactivity timeout: 1 hr Session timeout: 2 hr 30 min Session renewal: 30 min</pre> Web file transfer proxy: <pre>Proxy enabled: no</pre> Web file transfer certificate authority: <pre>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	
	3.3.5050	Updated Example
Role	admin	
Example	<pre> 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) # </pre>	

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	
	3.3.5050	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	
	3.3.5050	Added “ssl ciphers” parameter
		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</pre> Web User Interface: <pre>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.</pre> <pre>Inactivity timeout: 1 hr Session timeout: 2 hr 30 min Session renewal: 30 min</pre> Web file transfer proxy: <pre>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</pre> Web file transfer certificate authority: <pre>HTTPS server cert verify: yes HTTPS supplemental CA list: default-ca-list</pre> <pre>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
	3.3.5050 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	<pre>show web web proxy auth</pre>
Note	

4 System Management

4.1 Management Interface

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 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.

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

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

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

Step 3. Verify the IPv6 address is configured correctly.

```
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 10.209.0.2
switch (config)# show ip route
```

Destination	Mask	Gateway	Interface	Source
default	0.0.0.0	10.209.0.2	mgmt0	static
10.209.0.0	255.255.254.0	0.0.0.0	mgmt0	direct

```
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 system profiles 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:

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

Step 2. Create a VLAN interface. Run:

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

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

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

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

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

4.1.6 Commands

4.1.6.1 Interface

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

interface

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

Enters a management interface context.

Syntax Description	mgmt0	Management port 0 (out of band).
	mgmt1	Management port 1 (out of band).
	lo	Loopback interface.
	vlan<id>	In-band management interface (e.g. vlan10).
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # interface mgmt0 switch (config interface mgmt0) #</pre>	
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 TX bytes: 467577486 RX packets: 44866091 TX packets: 1385520 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	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	<pre>switch (config) # ip default-gateway mgmt1 switch (config) #</pre>	
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	<pre>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) #</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
 - 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	3.1.0000	
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: 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		

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 by using the “no” form of this command, or by enabling zeroconf.	

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 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 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 RX packets: 3738865 RX mcast packets: 0 RX discards: 0 RX errors: 0 RX overruns: 0 RX frame: 0 TX bytes: 40658219 TX packets: 142345 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	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	<code>ipv6 address</code> <code>show interface <ifname></code>
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	<pre>switch (config) # ipv6 dhcp primary-intf mgmt0 switch (config) #</pre>	
Related Commands	<pre>ipv6 enable ipv6 address show interface <ifname></pre>	
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	<pre>switch (config) # ipv6 dhcp stateless switch (config) #</pre>
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

show interface {<ifname> [configured | brief]}

Displays information about the specified interface, configuration status, and counters.

Syntax Description	ifname	The interface name e.g., “mgmt0”, “mgmt1”, “lo” (loopback), etc.
	configured	Displays the interface configuration.
	brief	Displays a brief info on the interface configuration and status.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.1.0000	
Role	admin	

Example

```
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:        my-interface
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) #
```

Related Commands	N/A

Notes

4.1.6.2 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	
Role	admin	
Example	<pre>switch (config) # hostname my-switch-hostname my-switch-hostname (config) #</pre>	
Related Commands	show hosts	
Notes	• Hostname may contain letters, numbers, and hyphens ('-'), in any combination • Hostname may not contain other letters, such as '%', '_', '.', etc • 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	

4.1.6.3 Routing

ip/ipv6 route

{ip | ipv6} route <network-prefix> <netmask> {<nexthop-address> | <ifname>}
no ip route <network-prefix> <netmask> {<nexthop-address> | <ifname>}

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	switch (config) # ip route 20.20.20.0 255.255.255.0 mgmt0				
	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
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 (IPv4 or 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) # ip default-gateway ::1 switch (config) # show ip 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 ip/ipv6 default-gateway

show {ip | 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) # ip default-gateway 10.10.10.10 switch (config) # show ip default-gateway Active default gateways: 172.30.0.1 (interface: mgmt0) switch (config) # show ip default-gateway static Configured default gateway: 10.10.10.10</pre>	
Related Commands	show ip default-gateway	
Notes	The configured IPv4 default gateway will not be used if DHCP is enabled.	

4.1.6.4 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	switch (config interface mgmt0) #ip arp 20.20.20.20 aa:aa:aa:aa:aa:aa			
	switch (config interface mgmt0) # show ip arp			
	Total number of entries: 6			
	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	switch (config interface mgmt0) #			
	show ip arp ip route			
Notes				

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: 60-28800.
Default	1500 seconds	
Configuration Mode	Config	
History	3.2.0230	
Role	admin	
Example	<pre>switch (config) # ip arp timeout 2000 switch (config) #</pre>	
Related Commands	<pre>ip arp show ip arp</pre>	
Notes	This value is used as the ARP timeout whenever a new IP interface is created.	

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	<pre>switch (config) # clear ipv6 neighbors switch (config) #</pre>
Related Commands	<pre>ipv6 neighbor show ipv6 neighbor arp</pre>
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 • See “clear ipv6 neighbors” on page 716 for the interface or VLAN specific command

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	<pre>switch (config) # show ipv6 neighbors IPv6 Address Age MAC Address State Interf ----- 2001::2 9428 AA:AA:AA:AA:AA:AA permanent mgmt0 switch (config) #</pre>	
Related Commands	<pre>ipv6 neighbor clear ipv6 neighbor show ipv6</pre>	
Notes		

4.1.6.5 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	<pre>switch (config) # ip dhcp default-gateway yield-to-static switch (config) # show ip dhcp DHCP DHCP Valid Interface Enabled Running 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) #</pre>	

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	

4.1.6.6 IP Diagnostic Tools

ping

ping [-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 http://linux.about.com/od/commands/l/blcmdl8_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.
	-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
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.

4.2.1 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: 04:21:44 Date: 2012/02/26 Time zone: America North United_States Other New_York 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> [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 are 3 or 4.
Default	NTP is enabled. NTP version number is 4.	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # no ntp peer 192.168.10.24 disable switch (config) #</pre>	
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.	

show clock

show clock

Displays the current system time, date and time zone.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
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	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # show ntp NTP is enabled. Clock is unsynchronized. No NTP peers or servers configured. switch (config) #</pre>
Related Commands	N/A
Notes	

•

4.3 Software Management

4.3.1 Upgrading MLNX-OS Software – Preconditions

Prior to upgrading MLNX-OS software from version 3.2.0100 and lower, please remove any old configuration from your system.

To remove old configuration:

Step 1. Clear your system of any old configuration. Run from CMM:

```
system:switch[2]> clear -cnfg
OK
system:switch[2]>
```

Step 2. Follow the steps described in [Section 4.3.2, “Upgrading MLNX-OS Software,”](#) on page 156.

4.3.2 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.



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.



To upgrade the MLNX-OS version of on a gateway, SM, or MLAG cluster, please refer to [Section 4.3.3, “Upgrading MLNX-OS HA Groups,”](#) on page 160.



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.



In order to upgrade the system on dual management system refer to [Section 4.3.2, “Upgrading MLNX-OS Software,”](#) on page 156.



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).

4.3.3 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.

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

Step 1. Identify the HA group master.

for MLAG. Run:

```
switch (config)# show mlag-vip
MLAG VIP
=====
MLAG group name: my-mlag-group
MLAG VIP address: 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
```

Step 2. Upgrade standby nodes in the HA group according to steps 1-8 in section [Section 4.3.2](#), on [page 156](#).

Step 3. Wait until all standby nodes have rejoined the group.

Step 4. Upgrade the master node in the HA group according to steps 1-8 in section [Section 4.3.2](#), on [page 156](#).

4.3.4 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-3.0.1224.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.

4.3.5 Downgrading MLNX-OS Software

IMPORTANT NOTE

If in possession of an SX65xx director switch with the notice presented in Figure 11, the lowest MLNX-OS version you can downgrade to is 3.3.5006; otherwise, the switch system will malfunction.

Figure 11: SX65xx Downgrade Attention Sticker



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.

4.3.5.1 Downloading Image

Step 1. Log into the system to obtain the serial number. Run:

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

Step 2. Download the requested MLNX-OS version from the following link:

<http://support.mellanox.com/SupportWeb/>

Step 3. Enter your username and password when prompted.

Step 4. Log into the switch via the CLI using the console port.

Step 5. Change to Config mode. Run:

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

Step 6. 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 7. Fetch the requested software image. Run:

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

4.3.5.2 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 partition (next). 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
```

4.3.5.3 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. Note that all configuration changes done with the new software are lost when returning to the older software version.

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

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

Also 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, in these cases above, 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
```

4.3.6 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.

4.3.6.1 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.

4.3.6.2 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.

Default Firmware Change on Standalone Systems

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

```
switch (config) # image fetch image fetch scp://root@1.1.1.1:/tmp/fw-SX-rel-9_2_6440-FIT.tgz
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.

4.3.7 Image Maintenance via Mellanox ONIE



Supported only on MSX1710-BS2F2O switch system.

The switch system MSX1710-BS2F2O allows booting ONIE and burning a different OS on the switch system.

When booting or rebooting the switch system an ONIE entry has been added to the boot loader options. For example:

```
GNU GRUB version 2.02~beta2

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
```



While MLNX-OS is installed, editing grub entry and grub command line are restricted.

ONIE may be selected from this prompt to allow ONIE functionality over this system. To do so, the MLNX-OS image burned must be uninstalled from the system. Once MLNX-OS is uninstalled, ONIE boots and the user is presented with ONIE command prompt which allows regular ONIE functionality according to Mellanox SwitchX ONIE Switch User Manual.

To return to MLNX-OS mode, MLNX-OS must be reinstalled using the ONIE Network OS installer file according to the preferred ONIE Network OS installation flow.

The switch system then loads from factory set configurations (automatically saved configuration is not supported).



All previous MLNX-OS installation flows are supported, therefore, the command “image fetch” or “image install” may be used to save previous configuration.

4.3.8 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	<pre>switch (config) # image boot location 2 switch (config) #</pre>	
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	<pre>switch (config) # boot next fallback-reboot enable switch (config) #</pre>
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	<pre>switch (config) # boot system location 2 switch (config) #</pre>	
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	<pre>switch (config) # image default-chip-fw image-SX_PPC_M460EX-ppc-m460ex- 20120122-084759.img switch (config) #</pre>	
Related Commands	<pre>image install-chip fw show images</pre>	
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]@host-name/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>	
Related Commands	show images	
Notes	• Please delete the previously available image, prior to fetching the new image • 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	<pre>switch (config) # image move image1.img image2.img switch (config) #</pre>	
Related Commands	show images	
Notes		

image options

image options {require-sig | serve}

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	Configures options for serving image files from this appliance
	all	Makes all image files on this appliance available for HTTP and HTTPS download
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # image options require-sig switch (config) #</pre>	
Related Commands	show images	
Notes		

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	

4.4 Configuration Management

4.4.1 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) #
```

4.4.2 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) #
```

4.4.3 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.

4.4.3.1 BIN Configuration Files

BIN configuration files are not human readable and cannot be edited.

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

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

➤ **To upload a BIN configuration file from a switch to an external file server**

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

➤ **To fetch a BIN configuration file**

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

➤ **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
```



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

4.4.3.2 Text Configuration Files

Text configuration files are text based and editable.

➤ **To create a new text-based configuration file:**

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

➤ **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 4.4.3.1](#)) 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. Reboot is not required.

4.4.4 Commands

4.4.4.1 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	<pre>switch (config) # debug generate dump Generated dump sysdump-switch-112104-201140526-091707.tgz switch (config) #</pre>
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	<pre>switch (config) # file debug-dump email sysdump-switch-112104-20114052-091707.tgz switch (config) #</pre>	
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]@host-name/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]@host-name/path/filename.
Default	N/A	
Configuration Mode	Config	
History	3.1.0000	
Role	admin	
Example	<pre>switch (config) # file tcpdump delete my-tcpdump-file.txt switch (config) #</pre>	
Related Commands	<pre>show files stats tcpdump</pre>	
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		

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	<pre>switch (config) # show files stats memory-201140524-111745.csv switch (config) #</pre>	
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	<pre>switch (config) # show files stats test dump3 switch (config) #</pre>
Related Commands	<pre>file tcpdump tcpdump</pre>
Notes	

4.4.4.2 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	<pre>switch (config) # configuration audit max-changes 100 switch (config) # show configuration audit Maximum number of changes to log: 100 switch (config) #</pre>	
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 file-name 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	<pre>switch (config) # configuration copy initial.bak example switch (config) #</pre>	
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]@host-name/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	<pre>switch (config) # configuration merge new-config-file switch (config) #</pre>	
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 switch-to

configuration switch-to <filename>

Loads the configuration from the specified file and makes it the active configuration file.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # show configuration files initial (active) newcon initial.prev initial.bak switch (config) # configuration switch-to newcon 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.

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	<pre>switch (config) # configuration text file my-config-file delete switch (config) #</pre>	
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	<pre>switch (config) # configuration write switch (config) #</pre>	
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	<pre>show running-config configuration write</pre>	
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	<pre>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) #</pre>	
Related Commands		
Notes		

show running-config

show running-config

Displays commands to recreate current running configuration.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config
History	3.1.0000
	3.3.4402 Removed “full” parameter
Role	monitor/admin
Example	<pre> switch (config) # show running-config ## ## Running database "initial" ## Generated at 2012/02/28 14:59:02 +0000 ## Hostname: switch-5ea5d8 ## ## ## License keys ## license install LK2-EFM_SX-5M11-5K11-5HGL-0KAL-64QK-8C2Q-60Q3-6C1G ## ## Network interface configuration ## interface mgmt0 create interface mgmt0 comment "" interface mgmt0 dhcp interface mgmt0 display interface mgmt0 duplex auto interface mgmt0 mtu 1500 no interface mgmt0 shutdown ... switch (config) # </pre>
Related Commands	
Notes	

4.5 Logging

4.5.1 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 18, “Supported Event Notifications and MIB Mapping,” on page 247](#).

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

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

Step 4. Override the log levels on a per-class basis. Run:

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

4.5.3 Switch Power-On Self-Test

As the switch powers on, it begins the Power On Self-Test (POST), a series of tests as part of its power-up procedure to ensure that the switch functions properly. During the POST, the switch logs any errors encountered. Some POST errors are critical, others are not.

The updated POST diagnostic code gets stored inside the “POST Diagnostic Register”.

Table 17 lists the POST return codes and their meanings.

Table 17 - POST Return Codes

Return Code	Severity	Meaning	POST Section
0x5	Critical	System initialization failure.	Standard POST

Table 17 - POST Return Codes

Return Code	Severity	Meaning	POST Section
0x10	Critical	Failure connecting to the main management process.	Standard POST
0x15	Critical	VPD initialization failure.	Standard POST
0x20	Critical	CPLD initialization failure.	Standard POST
0x25	Critical	Default IP configuration failure.	Standard POST
0x30	Critical	Temperature sensors failure.	Extended POST
0x35	Critical	Voltage sensors failure.	Extended POST
0x40	Critical	RAM memory failure.	Full POST
0x45	Critical	NAND memory failure.	Full POST
0x80	Non-critical	Incorrect firmware version.	Standard POST
0xff	Non-critical	POST ended successfully	Standard POST

4.5.4 Commands

logging <syslog IP address>

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
```

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]@hostname/path/filename)
Default	N/A	
Configuration Mode	Config	
History	3.3.4150	
Role	admin	
Example	<pre>switch (config) # logging debug-files delete current switch (config) #</pre>	
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

```
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) #
```

Related Commands

```
show logging
logging local
```

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	<pre>switch (config) # logging files delete current switch (config) #</pre>	
Related Commands	<pre>show logging show log files</pre>	
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/file-name.
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 uplaod 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 • warning • 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
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	severity 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.	

4.6 Debugging

➤ *To use the debugging logs feature:*

Step 1. Enable debugging. Run:

Step 2. Display the debug level set. Run:

Step 3. Display the logs. Run:

4.6.1 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	<pre>switch (config) # debug ethernet all switch (config) #</pre>
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	<pre>switch (config) # debug ethernet ip all switch (config) #</pre>
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	<pre>switch (config) # debug ethernet ip arp all switch (config) #</pre>
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	<pre>switch (config) # debug ethernet ip arp all switch (config) #</pre>	
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	<pre>switch (config) # debug ethernet ip dhcp-relay all switch (config) #</pre>	
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	<pre>switch (config) # debug ethernet ip igmp-l3 all switch (config) #</pre>	
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	<pre>switch (config) # debug ethernet ip interface all switch (config) #</pre>	
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	<pre>switch (config) # debug ethernet ip ospf all switch (config) #</pre>	
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	<pre>switch (config) # debug ethernet lacp all switch (config) #</pre>	
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	<pre>switch (config) # debug ethernet port all switch (config) #</pre>
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	<pre>switch (config) # debug ethernet port all switch (config) #</pre>	
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	<pre>switch (config) # debug ethernet spanning-tree all switch (config) #</pre>	

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	<pre>switch (config) # debug ethernet vlan all switch (config) #</pre>	
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: >>QoSQueueDelete 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: >>QoSQueueDelete 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], ulConfigOption[6] Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS u4IfIndex[33], ulConfigOption[6] Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS u4IfIndex[73], ulConfigOption[6] Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS u4IfIndex[121], ulConfigOption[6] Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS u4IfIndex[133], ulConfigOption[6] Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS u4IfIndex[13], ulConfigOption[6] Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS u4IfIndex[81], ulConfigOption[6] Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS u4IfIndex[117], ulConfigOption[6] Jun 15 16:20:49 switch-627e4c issd[6509]: TID 1274844336: [issd.DEBUG]: NPAPI: >>CfaGddConfigPort NS u4IfIndex[65], ulConfigOption[6] . . . switch (config) #</pre>	
Related Commands		
Notes		

4.7 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.

4.7.1 Supported Events

The following table presents the supported events and maps them to their relevant MIB OID.

Table 18 - 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	
disk-space-low	File system free space has fallen too low	Mellanox-EFM-MIB: diskSpaceLow	
health-module-status	Health module status changed	Mellanox-EFM-MIB: systemHealthStatus	
insufficient-fans	Insufficient amount of fans in system	Mellanox-EFM-MIB: insufficientFans	
insufficient-fans-recover	Insufficient amount of fans in system recovered	Mellanox-EFM-MIB: insufficientFansRecover	
insufficient-power	Insufficient power supply	Mellanox-EFM-MIB: insufficientPower	
interface-down	An interface's link state has changed to DOWN	RFC1213: linkdown (SNMPv1)	Supported for Ethernet, InfiniBand 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, InfiniBand and management interfaces for 1U and blade systems
internal-bus-error	Internal bus (I ² C) error	Mellanox-EFM-MIB: internalBusError	
liveness-failure	A process in the system is detected as hung	Not implemented	
low-power	Low power supply	Mellanox-EFM-MIB: lowPower	
low-power-recover	Low power supply recover	Mellanox-EFM-MIB: lowPowerRecover	

Table 18 - Supported Event Notifications and MIB Mapping

Event Name	Event Description	MIB OID	Comments
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 for SX65xx only systems
process-crash	A process in the system has crashed	Mellanox-EFM-MIB: procCrash	
process-exit	A process in the system unexpectedly exited	Mellanox-EFM-MIB: procUnexpectedExit	
snmp-authtrap	An SNMPv3 request has failed authentication	Not implemented	
topology_change	Topology change triggered by a local bridge	Bridge-MIB: topology-Change	Supported for Ethernet
unexpected-shutdown	Unexpected system shutdown	Mellanox-EFM-MIB: unexpectedShutdown	
To send, use the CLI command: <code>snmp-server notify send-test</code>	Send a testing event	testTrap	
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	

4.7.2 Terminal Notifications

➤ To print events to the terminal:

Set the events you wish to print to the terminal. Run:

```
switch (config) # logging monitor events notice
```

This command prints system events in the severity “notice” to the screen. For example, in case of interface-down event, the following gets printed to the screen.

```
switch (config) #
Wed Jul 10 11:30:42 2013: Interface IB1/17 changed state to DOWN
Wed Jul 10 11:30:43 2013: Interface IB1/18 changed state to DOWN
switch (config) #
```

4.7.3 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.

4.7.4 Commands

4.7.4.1 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 plaintext 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	<code>switch (config) # email autosupport ssl cert-verify</code>
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 “@<host-name>.<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	<pre>switch (config) # email autosupport enable switch (config) #</pre>
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	<code>switch (config) # email ssl cert-verify</code>
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	

4.8 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`.

4.8.1 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	<pre>switch (config) # no ha dns enable switch (config) #</pre>
Related Commands	
Notes	

4.9 User Management and Security

4.9.1 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 19 - 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`.

4.9.2 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.

4.9.2.1 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

4.9.2.2 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

4.9.2.3 LDAP

LDAP (Lightweight Directory Access Protocol) is an authentication protocol that allows a remote access server to forward a user's logon 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
```

4.9.3 Commands

4.9.3.1 User Accounts

username

username <username> [capability <cap> | disable [login | password] | 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
	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	
	3.3.5050	Updated Example
	3.4.1100	Updated Example
Role	admin	

Example

```

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) #

```

Related Commands

```

show usernames
show users

```

Notes

- To enable a user account, just set a password on it (or use the command `username <user> nopassword` to enable it with no password required for login)
- Removing a user account does not terminate any current sessions that user has open; it just prevents new sessions from being established
- Encrypted password is useful for the command `show configuration`, since the cleartext password cannot be recovered after it is set

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 System Administrator admin Password set 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	username show users
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) wtm 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	<pre>switch (config) # show whoami Current user: admin Capabilities: admin switch (config) #</pre>
Related Commands	<pre>username show usernames show users</pre>
Notes	

4.9.3.2 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 database 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 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 retroactive to previous logins. So if max-fail is disabled or changed, this does not immediately cause any users to be changed from locked to unlocked or vice-versa.
	unlock-time	Enables the auto-unlock of an account after a specified number of seconds if a user account is locked due to authentication failures, counting from the last valid login attempt. Unlike the “max-fail” setting, this does take effect immediately for all accounts. If both unlock-time and lock-time are set, the unlock-time must be greater than the lock-time. Careful with disabling the unlock-time, particularly if you have max-fail set to something, and have not overridden the behavior for the admin (i.e. they are subject to lockouts also). If the admin account gets locked out, and there are no other administrators who can aid, the user may be forced to boot single-user and use the <code>pam_tallybyname</code> command-line utility to unlock your account manually. Even if one is careful not to incur this many authentication failures, it makes the system more subject to DOS attacks.

Default	N/A
Configuration Mode	Config
History	3.2.3000
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	<pre>switch (config) # aaa authentication attempts reset user admin no-clear-history</pre>	
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	<pre>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) #</pre>	

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.
-
-

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	
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 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		

4.9.3.3 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	aaa authorization radius-server host show radius	
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	<pre>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) #</pre>	

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	aaa authorization radius-server radius-server host
Notes	

4.9.3.4 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	aaa authorization show radius show tacacs tacacs-server host	
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

```
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) #
```

Related Commands

```
aaa authorization
show tacacs
tacacs-server
```

Notes

- TACACS+ servers are tried in the order they are configured
- A PAP auth-type similar to an ASCII login, except that the username and password arrive at the network access server in a PAP protocol packet instead of being typed in by the user, so the user is not prompted
- If the user does not specify a parameter for this configured TACACS+ server, the configuration will be taken from the global TACACS+ server configuration. Refer to “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	

4.9.3.5 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	
	3.3.5050	Updated Example
Role	admin	
Example	<pre> 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) # </pre>	

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	
	3.3.5050	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	
	3.3.5050	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	
	3.3.5050	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	<pre>show aaa show ldap</pre>	
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 user-name.
	sAMAccountName	SAM Account name, active directory login name.
Default	sAMAccountName	
Configuration Mode	Config	
History	3.1.0000	
	3.3.5050	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	
	3.3.5050	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	show aaa show ldap	
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 3.3.5050 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	
	3.3.5050	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	show aaa show ldap	
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.
	3.3.5050	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	show aaa show ldap	
Notes	• If available, the TLS mode is recommended, as it is standardized, and may also be of higher security • The port number is used only for SSL mode. 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	
	3.3.5050	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	<pre> show aaa show ldap </pre>	
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	
	3.3.5050	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
	3.3.5050 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	<pre> show aaa show ldap </pre>
Notes	

4.10 Cryptographic (X.509, IPSec)

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.

4.10.1 Commands

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	<pre>switch (config)# crypto ipsec ike restart switch (config)#</pre>	
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** | **disable** | **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 Forwarding Secrecy) group to propose for Diffie-Hellman exchange for this IPsec peering • preshared-key: Configures the IKE pre-shared key for the IPsec peering • prompt-preshared-key: Prompts for the pre-shared key, rather than entering it on the command line • transform-set: Configures transform proposal parameters
	keying	Configures key management for this IPsec peering: • auth: Configures the authentication algorithm for this IPsec peering • disable: Configures this IPsec peering administratively disabled • encrypt: Configures the encryption algorithm for this IPsec peering • local-spi: Configures the local SPI for this manual IPsec peering • mode: Configures the peering mode for this IPsec peering • remote-spi: Configures the remote SPI for this manual IPsec peering
	manual	Configures IPsec peering using manual keys.

Default	N/A
Configuration Mode	Config
History	3.2.3000
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	<pre>switch (config) # crypto certificate name system-self-signed generate self-signed hash-algorithm sha256</pre>	
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	<pre>switch (config) # crypto certificate system-self-signed regenerate days-valid 3</pre>	
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 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	

4.11 Scheduled Jobs

Use the commands in this section to manage and schedule the execution of jobs

4.11.1 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	<pre>switch (config)# job 100 switch (config job 100) # enable switch (config job 100) #</pre>
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	<pre>switch (config)# job 100 switch (config job 100) # fail-continue switch (config job 100) #</pre>
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	<pre>switch (config)# job 100 switch (config job 100) # name my-job switch (config job 100) #</pre>	
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		

4.12 Statistics and Alarms

4.12.1 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: • <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) # stats alarm cpu_util_indiv clear switch (config) #</pre>	
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: • cpu_util_indiv - Average CPU utilization too high: percent utilization • disk_io - Operating System Disk I/O per second too high: kilobytes per second • fs_mnt - Free filesystem space too low: percent of disk space free • intf_util - Network utilization too high: bytes per second • memory_pct_used - Too much memory in use: percent of physical memory used • paging - Paging activity too high: page faults • temperature - Temperature is too high: degrees
Default	The default is different per alarm-id	
Configuration Mode	Config	
History	3.1.0000	
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	<pre>switch (config) # stats alarm cpu_util_indiv event-repeat single switch (config) #</pre>	
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	<pre>switch (config) # stats alarm cpu_util_indiv falling clear-threshold 10 switch (config) #</pre>	
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	<pre>switch (config) # stats alarm paging rate-limit window long 2000 switch (config) #</pre>	
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	<pre>switch (config) # stats chd memory_day clear switch (config) #</pre>	
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	<pre>switch (config) # stats chd memory_day enable switch (config) #</pre>	
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	<pre>switch (config) # stats sample temperature enable switch (config) #</pre>	
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	<pre>switch (config) # stats clear-all switch (config) #</pre>
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		

4.12.2 Power Management

4.12.2.1 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 361.

Table 20 - LWR Configuration Behavior

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

4.12.3 System Reboot

4.12.3.1 Rebooting 1U Switches

➤ *To reboot a 1U switch system:*

Step 1. Enter Config mode. Run:

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


Step 2. Reboot the system. Run:

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

4.12.4 Commands

4.12.4.1 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. Possible values: 10-2048.
	re-notif-cntr <counter>	Health control changes notification counter, in seconds. Possible values: 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		

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	<pre>switch (config) # power enable L01 switch (config) #</pre>	
Related Commands	<pre>show power show power consumers</pre>	
Notes	This command is not applicable for 1U systems.	

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 Ethernet	
History	3.3.4000	
Role	admin	
Example	<pre>switch (config interface ib 1/1) # power-management width auto switch (config) #</pre>	
Related Commands	show interface	
Notes		

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	<pre>switch (config) # usb eject switch (config) #</pre>
Related Commands	N/A
Notes	Applicable only for systems with USB interface.

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 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 SX_PPC_M460EX SX_3.4.0000 2014-10-14 20:26:41 ppc 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 output
Role	admin
Example	<pre>switch (config) # show uboot UBOOT version : U-Boot 2009.01 SX_PPC_M460EX SX_3.2.0330-82 ppc (Dec 20 2012 - 17:53:54) switch (config) #</pre>
Related Commands	N/A
Notes	

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 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
Role	admin
Example	<pre>switch (config) # show inventory ===== Module Type Part number Serial Number Asic revision ===== CHASSIS SX1036 MSX1036B-1SFR MT1205X01549 N/A MGMT SX1036 MSX1036B-1SFR MT1205X01549 0 FAN SXX0XX_FAN MSX60-PF MT1206X07209 N/A PS1 SXX0XX_PS MSX60-PF MT1206X06697 N/A 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 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 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
Role	admin
Example	<pre>switch (config) # show asic-version ===== Module Device Version ===== MGMT SX 9.2.9160 switch (config) #</pre>
Related Commands	N/A
Notes	

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																																																																													
Role	admin																																																																													
Example	<pre>switch (config) # show power</pre> <table><thead><tr><th>Module</th><th>Power (Watts)</th><th>Voltage</th><th>Current (Amp)</th><th>Capacity (Watts)</th><th>Grid Group</th><th>Status</th></tr></thead><tbody><tr><td>PS1</td><td>0.00</td><td>47.11</td><td>0.00</td><td>1008</td><td>A</td><td>OK</td></tr><tr><td>PS2</td><td>248.82</td><td>48.05</td><td>5.18</td><td>1008</td><td>A</td><td>OK</td></tr><tr><td>PS3</td><td>0.00</td><td>46.88</td><td>0.00</td><td>1008</td><td>A</td><td>OK</td></tr><tr><td>PS4</td><td>-</td><td>-</td><td>-</td><td>NOT PRESENT</td><td></td><td></td></tr><tr><td>PS5</td><td>46.72</td><td>47.82</td><td>0.98</td><td>1008</td><td>A</td><td>OK</td></tr><tr><td>PS6</td><td>-</td><td>-</td><td>-</td><td>NOT PRESENT</td><td></td><td></td></tr><tr><td>PS7</td><td>-</td><td>-</td><td>-</td><td>NOT PRESENT</td><td></td><td></td></tr><tr><td>PS8</td><td>-</td><td>-</td><td>-</td><td>NOT PRESENT</td><td></td><td></td></tr><tr><td>PS9</td><td>-</td><td>-</td><td>-</td><td>NOT PRESENT</td><td></td><td></td></tr><tr><td>PS10</td><td>-</td><td>-</td><td>-</td><td>NOT PRESENT</td><td></td><td></td></tr></tbody></table> <pre>Total power used : 295.54 W Total power capacity : 4032.00 W Total power budget : 4032.00 W Total power available : 3736.46 W Redundancy mode: combined Redundancy status: OK switch (config) #</pre>	Module	Power (Watts)	Voltage	Current (Amp)	Capacity (Watts)	Grid Group	Status	PS1	0.00	47.11	0.00	1008	A	OK	PS2	248.82	48.05	5.18	1008	A	OK	PS3	0.00	46.88	0.00	1008	A	OK	PS4	-	-	-	NOT PRESENT			PS5	46.72	47.82	0.98	1008	A	OK	PS6	-	-	-	NOT PRESENT			PS7	-	-	-	NOT PRESENT			PS8	-	-	-	NOT PRESENT			PS9	-	-	-	NOT PRESENT			PS10	-	-	-	NOT PRESENT		
Module	Power (Watts)	Voltage	Current (Amp)	Capacity (Watts)	Grid Group	Status																																																																								
PS1	0.00	47.11	0.00	1008	A	OK																																																																								
PS2	248.82	48.05	5.18	1008	A	OK																																																																								
PS3	0.00	46.88	0.00	1008	A	OK																																																																								
PS4	-	-	-	NOT PRESENT																																																																										
PS5	46.72	47.82	0.98	1008	A	OK																																																																								
PS6	-	-	-	NOT PRESENT																																																																										
PS7	-	-	-	NOT PRESENT																																																																										
PS8	-	-	-	NOT PRESENT																																																																										
PS9	-	-	-	NOT PRESENT																																																																										
PS10	-	-	-	NOT PRESENT																																																																										
Related Commands	N/A																																																																													
Notes																																																																														

show power consumers

show power consumers

Displays power consumers.

Syntax Description	N/A

Default	N/A

Configuration Mode	Any Command Mode

History	3.1.0000

Role	admin

Example	<pre>switch (config) # show power consumers ===== Module Power Voltage Current Status (Watts) ===== MGMT 17.47 48.00 0.36 OK S01 33.26 48.00 0.69 OK S02 33.50 48.00 0.70 OK L01 31.73 48.00 0.66 OK L02 29.76 48.00 0.62 OK L30 28.61 48.00 0.60 OK FAN5 14.91 48.00 0.31 OK FAN2 13.70 48.00 0.29 OK FAN1 14.21 48.00 0.30 OK FAN6 15.10 48.00 0.31 OK FAN4 14.53 48.00 0.30 OK FAN7 15.04 48.00 0.31 OK FAN3 15.17 48.00 0.32 OK FAN8 14.98 48.00 0.31 OK Total power used : 291.97 W Max power : 1636.00 W switch (config) #</pre>

Related Commands	N/A

Notes	

show temperature

show temperature

Displays the system's 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 voltage

show voltage

Displays power supplies voltage level.

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</pre> <pre>=====</pre> <table><thead><tr><th>Module</th><th>Power Meter</th><th>Reg</th><th>Expected Voltage</th><th>Actual Voltage</th><th>Status</th><th>High Range</th><th>Low Range</th></tr></thead><tbody><tr><td>MGMT</td><td>BOARD_MONITOR</td><td>USB 5V sensor</td><td>5.00</td><td>5.15</td><td>OK</td><td>5.55</td><td>4.45</td></tr><tr><td>MGMT</td><td>BOARD_MONITOR</td><td>Asic I/O sensor</td><td>2.27</td><td>2.11</td><td>OK</td><td>2.55</td><td>1.99</td></tr><tr><td>MGMT</td><td>BOARD_MONITOR</td><td>1.8V sensor</td><td>1.80</td><td>1.79</td><td>OK</td><td>2.03</td><td>1.57</td></tr><tr><td>MGMT</td><td>BOARD_MONITOR</td><td>SYS 3.3V sensor</td><td>3.30</td><td>3.28</td><td>OK</td><td>3.68</td><td>2.92</td></tr><tr><td>MGMT</td><td>BOARD_MONITOR</td><td>CPU 0.9V sensor</td><td>0.90</td><td>0.93</td><td>OK</td><td>1.04</td><td>0.76</td></tr><tr><td>MGMT</td><td>BOARD_MONITOR</td><td>1.2V sensor</td><td>1.20</td><td>1.19</td><td>OK</td><td>1.37</td><td>1.03</td></tr><tr><td>MGMT</td><td>CPU_BOARD_MONITOR</td><td>12V sensor</td><td>12.00</td><td>11.67</td><td>OK</td><td>13.25</td><td>10.75</td></tr><tr><td>MGMT</td><td>CPU_BOARD_MONITOR</td><td>12V sensor</td><td>2.50</td><td>2.46</td><td>OK</td><td>2.80</td><td>2.20</td></tr><tr><td>MGMT</td><td>CPU_BOARD_MONITOR</td><td>2.5V sensor</td><td>3.30</td><td>3.26</td><td>OK</td><td>3.68</td><td>2.92</td></tr><tr><td>MGMT</td><td>CPU_BOARD_MONITOR</td><td>SYS 3.3V sensor</td><td>3.30</td><td>3.24</td><td>OK</td><td>3.68</td><td>2.92</td></tr><tr><td>MGMT</td><td>CPU_BOARD_MONITOR</td><td>SYS 3.3V sensor</td><td>1.80</td><td>1.79</td><td>OK</td><td>2.03</td><td>1.57</td></tr><tr><td>MGMT</td><td>CPU_BOARD_MONITOR</td><td>1.8V sensor</td><td>1.20</td><td>1.24</td><td>OK</td><td>1.37</td><td>1.03</td></tr></tbody></table> <pre>switch (config) #</pre>								Module	Power Meter	Reg	Expected Voltage	Actual Voltage	Status	High Range	Low Range	MGMT	BOARD_MONITOR	USB 5V sensor	5.00	5.15	OK	5.55	4.45	MGMT	BOARD_MONITOR	Asic I/O sensor	2.27	2.11	OK	2.55	1.99	MGMT	BOARD_MONITOR	1.8V sensor	1.80	1.79	OK	2.03	1.57	MGMT	BOARD_MONITOR	SYS 3.3V sensor	3.30	3.28	OK	3.68	2.92	MGMT	BOARD_MONITOR	CPU 0.9V sensor	0.90	0.93	OK	1.04	0.76	MGMT	BOARD_MONITOR	1.2V sensor	1.20	1.19	OK	1.37	1.03	MGMT	CPU_BOARD_MONITOR	12V sensor	12.00	11.67	OK	13.25	10.75	MGMT	CPU_BOARD_MONITOR	12V sensor	2.50	2.46	OK	2.80	2.20	MGMT	CPU_BOARD_MONITOR	2.5V sensor	3.30	3.26	OK	3.68	2.92	MGMT	CPU_BOARD_MONITOR	SYS 3.3V sensor	3.30	3.24	OK	3.68	2.92	MGMT	CPU_BOARD_MONITOR	SYS 3.3V sensor	1.80	1.79	OK	2.03	1.57	MGMT	CPU_BOARD_MONITOR	1.8V sensor	1.20	1.24	OK	1.37	1.03
Module	Power Meter	Reg	Expected Voltage	Actual Voltage	Status	High Range	Low Range																																																																																																									
MGMT	BOARD_MONITOR	USB 5V sensor	5.00	5.15	OK	5.55	4.45																																																																																																									
MGMT	BOARD_MONITOR	Asic I/O sensor	2.27	2.11	OK	2.55	1.99																																																																																																									
MGMT	BOARD_MONITOR	1.8V sensor	1.80	1.79	OK	2.03	1.57																																																																																																									
MGMT	BOARD_MONITOR	SYS 3.3V sensor	3.30	3.28	OK	3.68	2.92																																																																																																									
MGMT	BOARD_MONITOR	CPU 0.9V sensor	0.90	0.93	OK	1.04	0.76																																																																																																									
MGMT	BOARD_MONITOR	1.2V sensor	1.20	1.19	OK	1.37	1.03																																																																																																									
MGMT	CPU_BOARD_MONITOR	12V sensor	12.00	11.67	OK	13.25	10.75																																																																																																									
MGMT	CPU_BOARD_MONITOR	12V sensor	2.50	2.46	OK	2.80	2.20																																																																																																									
MGMT	CPU_BOARD_MONITOR	2.5V sensor	3.30	3.26	OK	3.68	2.92																																																																																																									
MGMT	CPU_BOARD_MONITOR	SYS 3.3V sensor	3.30	3.24	OK	3.68	2.92																																																																																																									
MGMT	CPU_BOARD_MONITOR	SYS 3.3V sensor	1.80	1.79	OK	2.03	1.57																																																																																																									
MGMT	CPU_BOARD_MONITOR	1.8V sensor	1.20	1.24	OK	1.37	1.03																																																																																																									
Related Commands	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	

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 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 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
Role	admin
Example	<pre>switch (config) # show system capabilities IB: Supported Ethernet: Supported, Full L2 GW: Supported Max number of GW ports: 0 Max SM nodes: 648 IB Max licensed speed: FDR Ethernet Max licensed speed: 56Gb switch (config) #</pre>
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	<pre>switch (config) # show system mac 00:02:C9:5E:AF:18 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
	3.4.1100 Updated Example
Role	admin
Example	<pre>switch (config) # show protocols Ethernet enabled switch (config) #</pre>
Related Commands	N/A
Notes	

4.13 Network Management Interfaces

4.13.1 XML API

MLNX-OS XML API is currently under development. For further information please contact Mellanox support.

4.13.2 Commands

4.13.2.1 SNMP

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 [communities | mult-communities | notify]
no snmp-server enable [communities | mult-communities | notify]

Enables SNMP-related functionality.
 The no form of the command disables the SNMP server.

Syntax Description	enable	Enables SNMP-related functionality: • SNMP engine • SNMP traps
	communities	Enables community-based authentication on this system.
	mult-communities	Enables multiple communities to be configured.
	notify	Enables sending of SNMP traps and informs from this system.
Default	SNMP is enabled by default SNMP server communities are enabled by default SNMP notifies are enabled by default SNMP server multi-communities are disabled by default	
Configuration Mode	Config	
History	3.1.0000	First version
	3.2.1050	Change traps to notify
Role	admin	
Example	<pre>switch (config) # snmp-server enable 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 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 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.
 - If the command ends after the auth password, the privacy algorithm is set to its default, which is AES-128, and the privacy password is set to whatever was specified for the authentication password. You may also specify the privacy algorithm while still not specifying a separate password.
 - 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 switch(config) # </pre>
Related Commands	snmp-server auto-refresh
Notes	

4.13.2.2 XML API

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	<pre>switch (config) # xml-gw enable switch (config) # show xml-gw XML Gateway enabled: yes switch (config) #</pre>
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	Config
History	3.1.0000
Role	admin
Example	<pre>switch (config) # show xml-gw XML Gateway enabled: yes switch (config) #</pre>
Related Commands	xml-gw enable
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 – 1/10/40/56GbE (depends on the interface type and system)
- Description – user defined string
- Module-type – the type of the module plugged in the interface

5.1.1 Break-Out Cables

The break-out cable is a unique Mellanox capability, where a single physical 40Gbps port is divided into 2x10Gbps or 4x10Gbps ports. It maximizes the flexibility of the end user to use the Mellanox switch with a combination of 10Gbps and 40Gbps 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. 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 10G port (1,2 or 1,2,3,4). Each sub-physical port is then handled as an individual port. For example: splitting port 20 into 4 will give the following new ports: 1/20/1, 1/20/2, 1/20/3, 1/20/4.

Figure 12: Break-Out Cable



A split-4 operation results in blocking a 40G 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 40 Gb/s port into 4 lanes (4 SFP+ connectors). These 4 lanes then go, one lane to each of the 4 SFP+ connectors. Some ports can be split into 2 10 Gb/s ports, using lanes 1 and 2 only. When a QSFP port is split

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 to the table above

```
switch (config) # interface ethernet 1/19
switch (config interface ethernet 1/19) # shutdown
switch (config interface ethernet 1/19) # exit
switch (config) # interface ethernet 1/20
switch (config interface ethernet 1/20) # shutdown
```

Step 2. Split the ports as desired. Run:

```
switch (config interface ethernet 1/20) # module-type qsfp-split-4
switch (config interface ethernet 1/20) #
```

Step 3. The following warning will be displayed:

the following interfaces will be unmapped: 1/20 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 19 and 20 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 1/20/4) # shutdown
switch (config interface ethernet 1/20/4) # exit
switch (config) # interface ethernet 1/20/3
switch (config interface ethernet 1/20/3) # shutdown
switch (config interface ethernet 1/20/3) # exit
switch (config) # interface ethernet 1/20/2
switch (config interface ethernet 1/20/2) # shutdown
switch (config interface ethernet 1/20/2) # exit
switch (config) # interface ethernet 1/20/1
switch (config interface ethernet 1/20/1) # shutdown
```

Step 2. From the first member of the split (1/20/1), change the module-type back to QSFP. Run:

```
switch (config interface ethernet 1/20/1) # module-type qsfp
```



The module-type can be changed **only** from the first member of the split and **not** from the interface that was split.

The following warning will be displayed:

The following interfaces will be unmapped: 1/20/1 1/20/2 1/20/3 1/20/4.

Step 3. Type “yes” when prompted “Type 'yes' to confirm unsplit.”

5.1.2 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 1/20 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.

5.1.3 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		

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	

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	1522 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	<pre>switch (config interface ethernet 1/1) # mtu 9216 switch (config interface ethernet 1/1) #</pre>	
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	<pre>switch (config interface ethernet 1/1) # description my-interface switch (config interface ethernet 1/1) #</pre>	
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	1000 - 1GbE 10000 - 10GbE 40000 - 40GbE auto - auto negotiates link speed
	force	Forces speed change configuration
Default	Depends on the port module type, see the “Notes” section below.	
Configuration Mode	Config Interface Ethernet Config Interface MLAG Port Channel	
History	3.1.0000	
	3.3.3500	Added auto negotiation option
	3.3.4500	Added MLAG port-channel configuration mode
Role	admin	
Example	<pre>switch (config interface ethernet 1/1) # speed 40000 switch (config interface ethernet 1/1) #</pre>	
Related Commands	show interfaces ethernet	
Note	56Gbps port speed requires a license (LIC-6036F-56GE) - The default speed depends on the interface capabilities, interface capable with 40Gbps will have 40Gbps speed by default - 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	<pre>switch (config interface ethernet 1/1) # ip address dhcp switch (config interface ethernet 1/1) #</pre>
Related Commands	interface ethernet show interfaces ethernet
Note	

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> [counters [priority]]

Displays the configuration and status for the interface.

Syntax Description	inf	Interface number: <slot>/<port>.
	counters	Displays interface extended counters.
	priority	Displays interface extended counters per priority (0-7).
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.1.0000	
Role	admin	
Example	<pre> switch (config) # show interfaces ethernet 1/1 Eth1/1 Admin state: Enabled Operational state: Up Description: N/A Mac address: 00:02:c9:71:ed:2d MTU: 1500 bytes (Maximum packet size 1522 bytes) Flow-control: receive off send off Actual speed: 40 Gbps Width reduction mode: Not supported Switchport mode: access Last clearing of "show interface" counters 00:20:39 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 63 packets 0 unicast packets 63 multicast packets 0 broadcast packets 4032 bytes 0 discard packets switch (config) # </pre>	
Related Commands		
Note		

show interfaces ethernet [<inf>] capabilities

show interfaces ethernet [<inf>] capabilities

Displays the interface capabilities.

Syntax Description	inf	shows only one interface capabilities. 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 [<inf>] 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 Eth1/60 Admin state: Enabled Operational state: Down (Suspend) switch (config) #</pre>	
Related Commands		
Note		

show interfaces ethernet [<inf>] 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			
	3.4.1100Updated 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 [<inf>] 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.

module-type

module-type <type> [force]

Splits or un-splits the interface, as desired.

Syntax Description	type	qsfp - Port runs at 40000/56000Mbps. qsfp-split-2 - Port is split and runs at 2X10000Mbps. qsfp-split-4 - Port is split and runs at 4X10000Mbps.
	force	force the split operation without asking for user confirmation.
Default	interface module type is qsfp (if the interface supports 40Gbps speed)	
Configuration Mode	Config Interface Ethernet	
History	3.1.1400	
Role	admin	
Example	<pre>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) #</pre>	
Related Commands	switchport mode switchport [trunk hybrid] allowed-vlan show vlan	
Note	- The affected interfaces should be disabled prior to the operation - in order to un-split the interface - used the command with “qsfp”, the speed is set to 40Gbps “module-type qsfp”. - This command is applicable only on 40Gbps Ethernet ports	

5.2 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.

5.2.1 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.

5.2.2 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) #
```


5.2.3 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	<pre>switch (config interface port channel 10) # ip address dhcp switch (config interface port channel 10) #</pre>
Related Commands	<pre>interface port-channel show interface port-channel</pre>
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 LACPDUs Port Sent Recv Sent Recv 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 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	<pre>switch (config) # show lacp Port-channel Module Admin Status is enabled switch (config) #</pre>
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	<pre>switch (config)# show lacp interfaces port-channel 2 system-identifier Priority: 12345 MAC: 00:02:C9:AC:2A:60 switch (config)#</pre>	
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	Update Example
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: Not supported 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 4 discard 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	

5.3 MLAG



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 supported on Ethernet blades' internal as well as external ports.



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 be constructed of either 10GbE or 40GbE links. In such a case, PFC must be configured on this IPL. [Figure 15, “Basic MLAG Topology,” on page 442](#) illustrates. 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 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 6.7, “VRRP,” on page 902](#) or [Section 6.8, “MAGP,” on page 917](#) 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*].

5.3.1 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.

5.3.2 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.

5.3.3 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.

5.3.4 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.

5.3.5 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 457](#)) before closing its ports.

For more information on MLAG upgrade, please see [Section 4.3.3, “Upgrading MLNX-OS HA Groups,”](#) on [page 160](#).

5.3.6 MLAG Configuration

This section provides an example of how to configure two switches and a server in an MLAG setup.



➤ 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)# ipl 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 bpdufilter 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. or

➤ **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)
-----
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)#

```

5.3.7 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	<pre>switch (config) # protocol mlag switch (config) #</pre>
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, multi-cast 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

Enables MLAG.
 The no form of the command disables MLAG.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config MLAG
History	3.3.4500
Role	admin
Example	<pre>switch (config mlag) # no shutdown switch (config mlag) #</pre>
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	<pre>switch (config) # interface mlag-port-channel 1 switch (config interface mlag-port-channel 1) #</pre>	
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	<pre>switch (config interface vlan 1)# ipl 1 peer-address 10.10.10.10 switch (config interface vlan 1)#</pre>	
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	<pre>switch (config mlag) # keep-alive-interval 1 switch (config mlag) #</pre>	
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	<pre>switch (config interface ethernet 1/1)# mlag-channel-group 1 mode on switch (config interface ethernet 1/1)#</pre>	
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	<pre>switch (config)# mlag-vip my-mlag-domain ip 10.10.10.254/24 switch (config)#</pre>	
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	<pre>switch (config mlag) # reload-delay 30 switch (config mlag) #</pre>	
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	<pre>switch (config mlag) # system-mac 00:00:5E:00:01:5D switch (config mlag) #</pre>	
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	<pre>switch (config mlag) # upgrade-timeout 60 switch (config mlag) #</pre>	
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</pre> <table><thead><tr><th>Hostname</th><th>VIP-State</th><th>IP Address</th></tr></thead><tbody><tr><td>SwitchA</td><td>master</td><td>10.10.10.1</td></tr><tr><td>SwitchB</td><td>standby</td><td>10.10.10.2</td></tr></tbody></table> <pre>switch (config)#</pre>			Hostname	VIP-State	IP Address	SwitchA	master	10.10.10.1	SwitchB	standby	10.10.10.2
Hostname	VIP-State	IP Address										
SwitchA	master	10.10.10.1										
SwitchB	standby	10.10.10.2										
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
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: Not supported 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 discard 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	<pre>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) ----- 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)#</pre>		
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	

5.4 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.

5.4.1 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) #
```

5.4.2 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 hybrid 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) #
```

5.4.3 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) #
```

5.4.4 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) #
```

5.4.5 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> <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></td><td></td></tr> </tbody> </table> <pre>switch (config vlan 10) #</pre>		VLAN	Name	Ports	1	default	Eth1/2, Eth1/3, Eth1/4/1, Eth1/4/2 ...	10		
VLAN	Name	Ports									
1	default	Eth1/2, Eth1/3, Eth1/4/1, Eth1/4/2 ...									
10											
Related Commands	show vlan switchport mode switchport [trunk hybrid] allowed-vlan										
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></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></td><td></td><td>Eth1/16, Eth1/17, Eth1/18, Eth1/19, Eth1/20,</td></tr> <tr> <td></td><td></td><td>Eth1/21, Eth1/22, Eth1/23, Eth1/24, Eth1/25,</td></tr> <tr> <td></td><td></td><td>Eth1/26, Eth1/27, Eth1/28, Eth1/29, Eth1/30,</td></tr> <tr> <td></td><td></td><td>Eth1/31, Eth1/32, Eth1/33, Eth1/34, Eth1/35,</td></tr> <tr> <td></td><td></td><td>Eth1/36, Po34, Po4096</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, Eth1/5,			Eth1/6, Eth1/7, Eth1/8, Eth1/9, Eth1/10, Eth1/11, Eth1/12, Eth1/13, Eth1/14, Eth1/15,			Eth1/16, Eth1/17, Eth1/18, Eth1/19, Eth1/20,			Eth1/21, Eth1/22, Eth1/23, Eth1/24, Eth1/25,			Eth1/26, Eth1/27, Eth1/28, Eth1/29, Eth1/30,			Eth1/31, Eth1/32, Eth1/33, Eth1/34, Eth1/35,			Eth1/36, Po34, Po4096	10	my-vlan-name	
VLAN	Name	Ports																											
1	default	Eth1/2, Eth1/3, Eth1/4/1, Eth1/4/2, Eth1/5,																											
		Eth1/6, Eth1/7, Eth1/8, Eth1/9, Eth1/10, Eth1/11, Eth1/12, Eth1/13, Eth1/14, Eth1/15,																											
		Eth1/16, Eth1/17, Eth1/18, Eth1/19, Eth1/20,																											
		Eth1/21, Eth1/22, Eth1/23, Eth1/24, Eth1/25,																											
		Eth1/26, Eth1/27, Eth1/28, Eth1/29, Eth1/30,																											
		Eth1/31, Eth1/32, Eth1/33, Eth1/34, Eth1/35,																											
		Eth1/36, Po34, Po4096																											
10	my-vlan-name																												
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	<pre>switch (config interface ethernet 1/1) # switchport dot1q-tunnel qos- mode uniform switch (config interface ethernet 1/1) #</pre>	
Related Commands	<pre>show vlan show interfaces switchport switchport access vlan switchport [trunk hybrid] allowed-vlan vlan</pre>	
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-id	1-4094.																																								
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 expect this VLAN or VLAN range.
	none	Removes all VLANs.
Default	N/A	
Configuration Mode	Config Interface Ethernet Config Interface Port Channel Config Interface MLAG Port Channel	
History	3.1.1400	
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	<pre>show vlan show interfaces switchport switchport access vlan switchport mode vlan</pre>	
Note	This command is not applicable for interfaces with port mode access or access-dcb.	

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	

5.5 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.

5.5.1 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.

5.5.2 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: Not supported
  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) #
```

5.5.3 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	<pre>switch (config interface ethernet 1/1) # switchport dot1q-tunnel qos- mode uniform switch (config interface ethernet 1/1) #</pre>	
Related Commands	<pre>show vlan show interfaces switchport switchport access vlan switchport [trunk hybrid] allowed-vlan vlan</pre>	
Note		

5.6 MAC Address Table

5.6.1 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 unicast 00:11:22:33:44:55 vlan 1 interface
ethernet 0/1
```

5.6.2 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 16: MAC Learning Disable Example Case



5.6.3 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	<pre>switch (config) # clear mac-address-table dynamic switch (config) #</pre>
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	<pre>mac-address-table static clear mac-address-table</pre>	
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 Mac Address Aging Time: 300 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.

5.7 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 BPDUs that will have been received on the port.

5.7.1 Port Priority and Cost

When two ports on a switch are part of a loop, the STP port priority and port path cost configuration determine which port on the switch is put in the forwarding state and which port is put in the blocking state.

To configure port priority use the following command:

```
switch (config interface ethernet <inf>)# spanning-tree port-priority <0-240>
```

To configure port path cost use the following command:

```
switch (config interface ethernet <inf>)# spanning-tree cost <1-200000000>
```

5.7.2 Port Type

Port type has the following configuration options:

- **edge** – is not assumed to be converged by the RSTP learning/forwarding mechanism. It converges to forwarding quickly.



It is recommended to configure the port type for all ports connected to hosts as edge ports.

- **normal** – is assumed to be connected to a switch, thus it tries to be converged by the RSTP learning/forwarding. However, if it does not receive any BPDUs, it is operationally moved to be edge.
- **network** – is assumed to be connected 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}
```

5.7.3 BPDU Filter

Using BPDU filter prevents the CPU from sending/receiving BPDUs on specific ports.

BPDU filtering is configured per interface. When configured, the port does not send any BPDUs and drops all BPDUs that it receives. To configure BPDU filter, use the following command:

```
switch (config interface ethernet <inf>)# spanning-tree bpdupfilter {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.

5.7.4 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
```

5.7.5 Loop Guard

Loop guard is a feature that prevents loops in the network.

When a blocking port in a redundant topology transitions to the forwarding state (accidentally), an STP loop occurs. This happens when BPDUs are no longer received by one of the ports in a physically redundant topology.

Loop guard is useful in switched networks where devices are connected point-to-point. A designated bridge cannot disappear unless it sends an inferior BPDU or brings the link down on a point-to-point connection.



The loop guard configuration is only allowed on “network” 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
```

5.7.6 Root Guard

Configuring root guard on a port prevents that port from becoming a root port. A port put in root-inconsistent (blocked) state if an STP convergence is triggered by a BPDU that makes that port a root port. The port is unblocked after the port stops sending BPDUs.

To configure loop guard use the following command:

```
switch (config interface ethernet <inf>)# spanning-tree guard root
```

5.7.7 MSTP

Spanning Tree Protocol (STP) is a mandatory protocol to run on L2 Ethernet networks to eliminate network loops and the resulting broadcast storm caused by these loops. Multiple STP (MSTP) enables the virtualization of the L2 domain into several VLANs, each governed by a separate instance of a spanning tree which results in a network with higher utilization of physical links while still keeping the loop free topology on a logical level.

Up to 64 MSTP instances can be defined in the switch. Up to 64 VLANs can be mapped to a single MSTP instance. MSTP instance 0 (the default instance) may have all possible VLANs (1-4094) mapped to it.


For MSTP network design over Mellanox L2 VMS, please refer to [Mellanox Virtual Modular Switch Reference Guide](#).

5.7.8 RPVST

Rapid Per-VLAN Spanning Tree (RPVST) flavor of the STP provides finer-grained traffic by paving a spanning-tree instance per each configured VLAN. Like MSTP, it allows a better utilization of the network links comparing to RSTP.

Figure 17 exhibits a typical RPVST network configuration to get a better utilization on the inter-switch trunk ports.

Figure 17: RPVST Network Config



5.7.8.1 RPVST and VLAN Limitations

When the STP of the switch is set to RPVST, spanning tree is set on each of the configured VLANs in the system by default. To enable the spanning tree mode, the command “spanning-tree” must be run.

Each VLAN runs an STP state machine and an RPVST instance. There is a global limitation on the number of active state machines that can operate in MLNX-OS. Enforcement of this limitation is done through the maximum number of VLANs allowed in the system.

Table 23 - Supported VLANs by RPVST per Switch System

Switch System Model	Number of Supported VLANs
SX90Y3474	17


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.

5.7.8.2 RPVST and RSTP Interoperability

Figure 18: 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.

5.7.9 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	<pre>switch (config) # spanning-tree port type edge default switch (config) #</pre>	
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	<pre>switch (config) # spanning-tree priority 4096 switch (config) #</pre>	
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	<pre>switch (config) # interface ethernet 1/1 switch (config interface ethernet 1/1) # spanning-tree port-priority 16 switch (config interface ethernet 1/1) #</pre>	
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 0-200000000.
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	<pre>switch (config) # interface ethernet 1/1 switch (config interface ethernet 1/1) # spanning-tree cost 1000 switch (config interface ethernet 1/1) #</pre>	
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	<pre>switch (config) # interface ethernet 1/1 switch (config interface ethernet 1/1) # spanning-tree port type edge switch (config interface ethernet 1/1) #</pre>	
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	<pre>switch (config) # interface ethernet 1/1 switch (config interface ethernet 1/1) # spanning-tree guard root switch (config interface ethernet 1/1) #</pre>	
Related Commands	show spanning-tree	
Note		

spanning-tree bpdudfilter

spanning-tree bpdudfilter {disable | enable}
no spanning-tree bpdudfilter

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	<pre>switch (config) # interface ethernet 1/1 switch (config interface ethernet 1/1) # spanning-tree bpdudfilter enable</pre>	
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	<pre>switch (config) # clear panning-tree counters switch (config) #</pre>
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	<pre>switch (config)# spanning-tree mst max-hops 20 switch (config)#</pre>	
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	<pre>switch (config)# spanning-tree mst 1 vlan 10-20 switch (config)#</pre>	
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	<pre>switch (config)# spanning-tree mst revision 1 switch (config)#</pre>	
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	<pre>switch (config)# spanning-tree mst name my-mst switch (config)#</pre>	
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	<pre>switch (config)# spanning-tree mst name my-mst switch (config)#</pre>	
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	<pre>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</pre>	
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 0-2000000000.
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	clear spanning-tree counters spanning-tree					
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	<pre>switch (config) # show spanning-tree root</pre> <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></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> <pre>switch (config) #</pre>								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	<pre> clear spanning-tree counters spanning-tree </pre>	
Note		

5.8 OpenFlow

MLNX-OS supports OpenFlow 1.0. 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 datapath 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.

5.8.1 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 following 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:
 - a. Forwarding the packet
 - b. Dropping the packet
 - c. Forwarding the packet to the OpenFlow controller
 - d. Modifying the VLAN, VLAN priority (PCP), and/or stripping the VLAN header



The flow table supports up to 1000 flows.

5.8.2 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
```

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).

5.8.3 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	<pre>switch (config) # protocol openflow switch (config) #</pre>
Related Commands	
Note	

openflow description

openflow description <string>

Sets the OpenFlow description.

Syntax Description	string	Free string.
Default	N/A	
Configuration Mode	Config	
History	3.3.4302	
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		

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
Role	admin
Example	switch (config interface ethernet 1/1)# openflow mode hybrid switch (config interface ethernet 1/1)#
Related Commands	
Note	

controller-ip

controller-ip <ip-address> [tcp-port <port-number>]
no controller-ip <ip-address> 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 <port-number>	Sets the TCP port number of the OpenFlow controller.
Default	0.0.0.0; TCP port 6633	
Configuration Mode	Config OpenFlow	
History	3.3.4200	
Role	admin	
Example	<pre>switch (config openflow) # controller-ip 10.10.10.10 tcp-port 6633 switch (config openflow) #</pre>	
Related Commands		
Note		

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	<pre>switch (config openflow) # datapath-id 0x1234 switch (config openflow) #</pre>	
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		

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	<pre>switch (config openflow) # show openflow flows 2 Flow id: 2 priority: 1 hard timeout: infinite idle timeout: 0 sec match: ingress interface: Eth1/18 source Ethernet address: 11:22:33:44:55:66 destination Ethernet address: 77:88:aa:bb:cc:fe Ethernet type: 0x800 VLAN ID: 308 PCP: 4 SIP: 1.1.1.1 DIP: 2.2.2.2 Dport: 1790 Protocol: 86 TOS: 120 actions: output controller statistics: 0 packets, 0 bytes switch (config openflow) #</pre>	
Related Commands		
Note		

5.9 IGMP Snooping



Only IGMP Snooping v1 and v2 are supported.

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 router-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 router-port, and the router-port updates its multicast-group data-base accordingly. Each dynamically learnt multicast group will be added to all of the router-ports on the switch.

As many as 5K multicast groups can be created on the switch.

5.9.1 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
```

5.9.2 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) #
```

5.9.3 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) # 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
VLAN 10  IGMP Querier Present  query-interval: 125  address: 1.1.1.2  version: 2
switch (config vlan 10) #
```

5.9.4 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	<pre>switch (config) # ip igmp snooping switch (config) # vlan 10 switch (config vlan 10) # ip igmp snooping switch (config vlan 10) # exit switch (config) # show ip igmp snooping</pre> IGMP snooping global configuration: <pre>IGMP snooping globally enabled 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</pre> <pre>switch (config vlan 10) # show ip igmp snooping vlan 10</pre> Vlan 10 configuration parameters: <pre>IGMP snooping is enabled IGMP version is V2 Snooping switch is acting as Non-Querier mrouter static port list: none mrouter dynamic port list: none switch (config vlan 10) # vlan 10 switch (config vlan 10) #</pre>
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 router 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 router 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
switch (config) # show ip igmp snooping
```

```
IGMP snooping global configuration:
```

```
IGMP snooping globally enabled
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 3 seconds
```

```
switch (config) #
```

Related Commands

```
ip igmp snooping (admin)
show ip igmp snooping
```

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	<pre> switch (config interface ethernet 1/1) # ip igmp snooping fast-leave switch (config interface ethernet 1/1) # show ip igmp snooping inter- faces interface leave-mode ----- Eth1/1 Fast Eth1/2 Normal Eth1/3 Normal ... switch (config interface ethernet 1/1) # </pre>
Related Commands	show ip igmp snooping interfaces
Note	

ip igmp snooping static-group

ip igmp snooping static-group <IP address> interface <type> <number>
no ip igmp snooping static-group <IP address> interface <type> <number>

Creates a static multicast group and attaches a port to a specified group.
 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 <type> <number>	Attach the group to a specific interface. type - ethernet or port-channel	
Default	No static groups are configured.		
Configuration Mode	Config VLAN		
History	3.1.1400		
Role	admin		
Example	<pre>switch (config)# vlan 1 switch (config vlan 1) # ip igmp snooping static-group 230.0.0.1 interface ethernet 1/1 switch (config vlan 1) # show ip igmp snooping groups Vlan ID Group St/Dyn Ports ----- 1 230.0.0.1 St Eth1/1 Total Num of Dynamic Group Addresses 0 Total Num of Static Group Addresses 1 switch (config vlan 1) #</pre>		
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 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.	Attaches the group to a specific interface.
Default	No static mrouter are configured.	
Configuration Mode	Config VLAN	
History	3.1.1400	
Role	admin	
Example	<pre>switch (config)# vlan 1 switch (config vlan 1) # ip igmp snooping mrouter interface ethernet 1/1 switch (config vlan 1) # show ip igmp snooping mrouter Vlan Ports ----- 1 Eth1/1(static) switch (config vlan 1) #</pre>	
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 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	<pre>switch (config) # ip igmp snooping unregistered multicast flood switch (config) # show ip igmp snooping</pre> IGMP snooping global configuration: <pre>IGMP snooping globally enabled 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	show ip igmp snooping
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	<pre>switch (config vlan 1)# ip igmp snooping querier switch (config vlan 1)#</pre>
Related Commands	<pre>igmp snooping querier query-interval show ip igmp snooping querier</pre>
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	<pre>switch (config vlan 1)# igmp snooping querier query-interval 20 switch (config vlan 1)#</pre>	
Related Commands	<pre>igmp snooping querier query-interval show ip igmp snooping querier</pre>	
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
Role	admin
Example	<pre>switch (config) # show ip igmp snooping IGMP snooping global configuration: IGMP snooping globally enabled 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 3 seconds IGMP snooping unregistered multicast: flood switch (config) #</pre>
Related Commands	
Note	

show ip igmp snooping groups

show ip igmp snooping groups

Displays per VLAN the list of multicast groups attached (static or dynamic allocated) per port.

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 groups Vlan ID Group St/Dyn Ports ----- 1 230.0.0.1 St Eth1/1 Total Num of Dynamic Group Addresses 0 Total Num of Static Group Addresses 1 switch (config) #</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		

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 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 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
Role	admin
Example	<pre>switch (config) # show ip igmp snooping statistics Snooping Statistics for VLAN 1 General queries received : 0 Group specific queries received : 0 V1/V2 reports received : 0 V1/V2 reports transmitted : 0 Leave messages received : 0 Group specific queries transmitted: 0 Leave messages transmitted: 0 Unsuccessful joins received count Per Vlan: 0 Active/Successful joins received count Per Vlan: 0 Active Groups count: 0 Packets dropped: 0 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	
Role	admin	
Example	<pre>switch (config) # show ip igmp snooping querier vlan 10 Vlan 1 IGMP Querier Present query-interval: 20 address: 1.1.1.2 version: 2 switch (config) #</pre>	
Related Commands		
Note		

5.10 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.

5.10.1 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
```

5.10.2 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.

5.10.3 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	<pre>switch (config)# lldp tx-delay 10 switch (config)#</pre>	
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	<pre>switch (config)# lldp tx-hold-multiplier 10 switch (config)#</pre>	
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 {receive | transmit}

lldp {receive | transmit}
no lldp {receive | transmit}

Enables LLDP to be received or transmitted on this port.
 The no form of the command disables the LLDP to be received or transmitted on this port.

Syntax Description	N/A
Default	Enabled for receive and Transmit.
Configuration Mode	Config Interface Ethernet
History	3.2.0300
Role	admin
Example	<pre>switch (config interface ethernet 1/1)# lldp receive switch (config interface ethernet 1/1)# lldp transmit switch (config interface ethernet 1/1)#</pre>
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	<pre>switch (config interface ethernet 1/1)# lldp tlv-select port-description sys-name 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

Shows 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 interface

show lldp interface [ethernet <inf>]

Shows LLDP local interface table information.

Syntax Description	inf					Interface number (e.g. 1/1).				
Default	N/A									
Configuration Mode	Any Command Mode									
History	3.2.0300		First version							
	3.3.4200		Updated example							
	3.3.4402		Updated example							
Role	admin									
Example	switch (config)# show lldp interface ethernet 1/1									
	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									
Related Commands	Interface		Receive		Transmit		Notification		TLVs	

	eth1/1		Enable		Enable		Enable		PD, SN, SD, SC, MA, PFC, CEE	
	eth1/2		Disable		Disable		Enable		PD, SN, SD, MA, AP	
	eth1/3		Enable		Disable		Disable		PD, SD, SC, ETS-R, AP, PFC	
	...									
Note	switch (config)#									

show lldp interfaces ethernet <inf> remote

show lldp interfaces ethernet <inf> remote

Shows LLDP remote interface table information.

Syntax Description	inf	Local interface number (e.g. 1/1).
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.2.0300	First version
	3.3.4200	Updated output
Role	admin	

Example

```

switch (config)# show lldp interfaces ethernet <number> Ethernet <port-number> // example "Ethernet 1/1"
Latest LLDPDU received on <date> // e.g. date: "Thu Feb 14 12:08:29 2013" - new field
Remote Index:
Remote chassis id: <byte array> ; chassis id subtype: <sub-type>
Remote port-id: <byte array> ; port id subtype: <sub-type>
Remote port description: <byte array>
Remote system name: <byte array>
Remote system description: < byte array>
Remote system capabilities supported: <enum parsed as defined in the MIB> ; enable <enum parsed as defined in the
MIB>
Management Table //theoretially remote can send more then one management address (Future) RemoteIndex Subtype Address
ifSubtype ifId OID
1 ipV4(1) 10.10.10.10 ifIndex(2) 1(mgmt0) <OID>
1 ipV4(1) 10.10.10.11 ifIndex(3) 2(mgmt1) <OID>
Unknown TLVs Table //(Future) Type Info
-----
<integer> <byte-array>
<integer> <byte-array>
Organizationally-Defined Information Table // (Future) OUI subtype Index DefInfo
-----
<byte-array> <integer> <integer> <byte-array>
<byte-array> <integer> <integer> <byte-array>
Remote PFC configuration // new section Willing: {enabled, disabled}
MACsec: {enabled, disabled}
Number of supported traffic classes: 4 // range is 1-8
PFC enabled on priorities: 5 7 // it could be "0 1 2 3 4 5 6 7" or " 1 3 7" or "None"
WARNING: peer PFC configuration does not match the local PFC configuration // This warning should appear only if the
local and remote PFC configuration don't match!
Remote ETS configuration // new section Willing: {enabled, disabled}
CBS: {enabled, disabled}
Number of supported traffic classes: 3 // range is 1-8
WARNING: peer ETS configuration does not match the local ETS configuration // This warning should appear only if the
local and remote ETS configuration don't match!
Priority assignment table: Priority TC
-----
0
0
1
0
2
1
3
1
4
2
5
2
6
3
7
3
Traffic class bandwidth table // (No need to have the recommended TC - ETS-Recommended TLV)
TC Bandwidth TSA
-----
Page 43 of 354
0 25% tsaStrictPriority(0) // we are expecting either 0 or 2 as answers.

```

Related Commands**Note**

show lldp timers

show lldp timers

Shows 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 lldp statistics global

show lldp statistics global

Shows 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 statistics [interface ethernet <inf>]

show lldp statistics [interface ethernet <inf>]

Shows 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 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 Application priority configuration Selector Protocol Priority ----- Ethertype 0x8906 3 Ethertype 0x8914 3 switch (config)#</pre>
Related Commands	
Note	

5.11 Quality of Service (QoS)

5.11.1 Priority Flow Control and Link Level 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 has 8 possible priorities (3 bits in VLAN header). Each priority can be mapped to one of 4 possible queues in the ingress.

The PFC software 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 datapath
- Provides hardware abstraction layer callouts 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 datapath. 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
```

5.11.2 Enhanced Transmission Selection (ETS)

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) #
```

5.11.3 Commands

5.11.3.1 Enhanced Transmission Selection (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
Role	admin
Example	<pre>switch (config)# dcb ets enable switch (config)# show dcb ets</pre> ETS enabled <pre>TC Bandwidth ----- 0 25% 1 25% 2 25% 3 25%</pre> Number of Traffic Class: 4 <pre>switch (config) #</pre>
Related Commands	show dcb ets
Note	

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
Role	admin
Example	<pre>switch (config)# dcb ets tc bandwidth 20 20 30 30 switch (config) # show dcb ets</pre> ETS enabled <pre>TC Bandwidth ----- 0 20% 1 20% 2 30% 3 30%</pre> Number of Traffic Class: 4 <pre>switch (config) #</pre>
Related Commands	show dcb ets
Note	The sum of all traffic class bandwidth must be equal to 100.

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
Role	admin
Example	switch (config interface ethernet 1/1) # vlan map-priority 1 traffic-class 2 switch (config interface ethernet 1/1) #
Related Commands	show dcb ets interface
Note	

show dcb ets

show dcb ets

Displays ETS configuration and operational data.

Syntax Description

Default ETS is enabled.

Configuration Mode Any Command Mode

History 3.1.0000

Role admin

Example switch (config)# show dcb ets

ETS enabled

TC	Bandwidth
0	25%
1	25%
2	25%
3	25%

Number of Traffic Class: 4

switch (config) #

Related Commands

Note

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**

5.11.3.2 Priority Flow Control (PFC)

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	prio0-7.
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 3.3.4500 Added 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 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	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

```
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) #
```

Related Commands

Note

5.12 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.

5.12.1 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
```

5.12.2 ACL Actions

An ACL action is a set of actions can be activated in case the packet hits the ACL rule.

➤ **To modify the VLAN tag of the egress traffic as part of the ACL “permit” rule:**

Step 1. Create access-list action profile:

- a. Create an action access-list profile using the command `access-list action <action-profile-name>`
- b. Add rule to map a VLAN using the command `vlan-map <vlan-id>` within the action profile configuration mode

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 port access-group my-list
```

5.12.3 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	<pre>switch (config)# mac access-list my-mac-list switch (config mac access-list my-mac-list)#</pre>	
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:1-500.
	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
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:1-500.
	deny permit	Determines the type of the rule, deny or permit action. Valid mask values fall in the range 0-255.
	{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
Role	admin	
Example	<pre>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) #</pre>	
Related Commands	ipv4/mac access-list ipv4/mac port access-group	
Note		

deny/permit (IPv4 TCP/UDP ACL rule)

```
[seq-number <sequence-number>] {permit | deny} {tcp | udp} {<source-ip>
[mask <ip>] | [any]} {<dest-ip> [mask <ip>] | [any]} [eq-source <port-number>]
[eq-destination <port-number>] [action <action-id>]
no <sequence-number>
```

Creates a rule for IPv4 UDP/TCP ACL.

The no form of the command deletes a rule from the ACL.

Syntax Description	sequence-number	Optional parameter to set a specific sequence number for the rule. The range is:1-500.
	deny permit	Determines the type of the rule, deny or permit action.
	tcp udp	UDP or TCP 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 MAC. The “any” option will cause the rule not to check the destination MAC.
	[eq-source <port-number>]	TCP/UDP source port number. Range is 0-65535.
	[eq-destination <port-number>]	TCP/UDP destination port number. Range is 0-65535.
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	
Role	admin	
Example	<pre>switch (config ipv4 access-list my-list) # seq-number 10 deny tcp any any eq-source 1200 switch (config ipv4 access-list my-list) #</pre>	
Related Commands	<pre>ipv4/mac access-list ipv4/mac port access-group</pre>	
Note		

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> switch (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	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		

5.13 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 19 provides an overview of the mirroring functionality.

Figure 19: 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.

5.13.1 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 20: Mirror to Analyzer Mapping



A mirroring session is a monitoring configuration mode that has the following parameters:

Table 24 - 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

5.13.1.1 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.

5.13.1.2 Destination Interface

The destination interface is an analyzer port is one to which mirrored traffic is sent. The mirrored packets, are duplicated, optionally modified and sent to the analyzer port. The SwitchX® platform supports up to 7 analyzer ports where any mirror port can be mapped to any analyzer port and more than a single mirror port can be mapped to a single analyzer port.

Packets can be forwarded to any destination using the command `destination interface`.
The analyzer port supports status and statistics as any other port.

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).

5.13.1.3 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 21: Header Format Options

5.13.1.4 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 10G 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.

5.13.1.5 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.

5.13.2 Configuring Mirroring Sessions

Figure 22 presents two network scenarios with direct and remote connectivity to the analyzer equipment. Direct connectivity is when the analyzer is connected to the analyzer port of the switch. In this case there is no need for adding an L2 header to the mirrored traffic. Remote connectivity is when the analyzer is indirectly connected to the analyzer port of the switch. In this situation, adding an L2 header may be necessary depending on the network's setup.

Figure 22: 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
```

5.13.3 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), pol(e)
3	Enable	Up	add-eth	eth1/5	eth1/18(e)
7	Disable	Down	local		

5.13.4 Commands

5.13.4.1 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. The range is 1-7.
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		

5.13.4.2 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	<pre>switch (config monitor session 1) # destination interface ethernet 1/2 switch (config monitor session 1) #</pre>	
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	<pre>switch (config monitor session 1) # no shutdown switch (config monitor session 1) #</pre>
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	
Role	admin	
Example	<pre>switch (config monitor session 1) # add source interface ethernet 1/1 direction both switch (config monitor session 1) #</pre>	
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”.	

header-format

header-format {local [traffic-class <tc>] | add-vlan <vlan-id> [priority <prio>]
[traffic-class <tc>] | 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 is 0-3.
	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 is 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	
Role	admin	
Example	<pre>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)#</pre>	
Related Commands		
Note	If add-ethernet-header is used, the source MAC address is the one attached to the switch.	

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	<pre>switch (config monitor session 1) # truncate switch (config monitor session 1) #</pre>
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.	

5.13.4.3 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	
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		

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), pol(e) 3 Enable Up add-eth eth1/5 eth1/18(e) 7 Disable Down local switch (config) #</pre>
Related Commands	
Note	


5.14 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 23: sFlow Functionality Overview



5.14.1 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.

5.14.2 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 25 - List of Statistical Counters

Counter	Description
Total packets	The number of packets that pass through sFlow-enabled ports.

Table 25 - List of Statistical Counters

Counter	Description
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.

5.14.3 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.

5.14.4 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.

5.14.5 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
```

5.14.6 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

Interfaces
Ethernet: eth1/1
Port-channel: po1

Statistics:
Total Packets: 2000
Number of flow samples: 1200
Number of samples discarded: 0
Number of statistic samples: 800
Number of datagrams: 300
```

5.14.7 Commands

5.14.7.1 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	<pre>switch (config) # protocol sflow switch (config) #</pre>
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	

5.14.7.2 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	<pre>switch (config sflow) # sampling-rate 16111 switch (config sflow) #</pre>	
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	<pre>switch (config sflow) # max-sample-size 165 switch (config sflow) #</pre>	
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	<pre>switch (config sflow) # counter-poll-interval 30 switch (config sflow) #</pre>	
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	<pre>switch (config sflow) # agent-ip 20.20.20.20 switch (config sflow) #</pre>	
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	<pre>switch (config sflow) # clear counters switch (config sflow) #</pre>
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	

5.14.7.3 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
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 Interfaces Ethernet: eth1/1 Port-channel: po1 Statistics: Total Packets: 2000 Number of flow samples: 1200 Number of samples discarded: 0 Number of statistic samples: 800 Number of datagrams: 300</pre>
Related Commands	
Note	

5.15 Transport Applications

5.15.1 RDMA over Converged Ethernet (RoCE)


5.15.1.1 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 (InfiniBand) 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 InfiniBand 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.

5.15.1.1.1IP 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 24: 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 implement 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 25: RoCEv2 Protocol Stack



5.15.1.2 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>

5.15.1.2.1 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

1. Standard RDMA APIs are IP based already for all existing RDMA technologies

- All InfiniBand verbs applications which run over InfiniBand verbs should work on RoCE links if they use GRH headers.
- Set HCA to use Ethernet protocol:
Display the Device Manager and expand “System Devices”.

5.15.1.2.2 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.

5.15.1.2.2.1 Using Global Pause Flow Control (GFC)

➤ *To use Global Pause Flow Control (GFC) mode, disable QoS and Priority:*

```
PS $ Disable-NetQosFlowControl
PS $ Disable-NetAdapterQos
```

5.15.1.3 Configuring SwitchX® Based Switch System

➤ *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

For further information on how to configure SwitchX, please refer to SwitchX User Manual.

5.15.1.4 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.

5.15.1.4.1 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.

5.15.1.5 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 -RocMode 1
```

➤ ***To enable RoCEv2 using the PowerShell:***

- Open the PowerShell and run:

```
Set-MlnxDriverCoreSetting -RocMode 2
```

➤ ***To disable any version of RoCE using the PowerShell:***

Open the PowerShell and run:

```
Set-MlnxDriverCoreSetting -RocMode 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
.
.
.
RocMode          : 0
```

5.16 802.1x 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 ports.



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.

5.16.1 802.1x Operating Modes

The following operating modes are supported in 802.1x:

- Single host – only one supplicant can communicate through the port.
Once authentication of the supplicant is accepted by the authentication server, the switch allows it access. If the supplicant logs off or the port state is changed, the port becomes unauthenticated. And if a different supplicant tries to access through this port, its bidirectional traffic is discarded (including authentication traffic).



An exception to this is multicast and broadcast traffic which do get transmitted over the interface once authenticated and are exposed to an unauthorized supplicant if it exists.

- Multi-host mode – allows connection of multiple hosts over a single port. Only the first supplicant is authenticated. Subsequent hosts have network access without the need to authenticate.

5.16.2 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)#
```

5.16.3 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	<pre>switch (config interface ethernet 1/2)# dot1x timeout reauthentication 3600</pre>	
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	

6 IP Routing

6.1 General

6.1.1 IP Interfaces

MLNX-OS supports 3 types of IP interfaces.

- VLAN interface
- Loopback interface
- Router ports



Router ports 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 is a regular switch port configured to operate as an L3 interface. Router ports are assigned an IP address and all L3 commands become applicable to them.

Once configured, router ports 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

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 an L3 router port. Run:

```
switch (config interface ethernet 1/10)# no switchport force
```

Step 3. Configure an IP address on the router port. Run:

```
switch (config interface ethernet 1/10)# ip address 100.100.100.100 /24
```

Step 4. Verify router port 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: Unknown
  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 26, 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 26: 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 27: 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 IPv4 Routing Mode

The resources available for IPv4 routing are as follows:

- number of IPv4 neighbors – 2048
- number of IPv4 unicast routes – 4096
- number of IPv4 multicast routes – 672



Prior to upgrading to this software release the user must align the number of configured multicast routes to number defined above.

6.1.5 Commands

6.1.5.1 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
Configuration Mode	Config
History	3.4.1802
Role	admin
Example	<pre>switch (config) # ip l3 force switch (config) #</pre>
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	<pre>switch (config) # vrf definition my-vrf switch (config vrf definition my-vrf) #</pre>	
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	<pre>switch (config) # routing-context vrf my-vrf switch (config) #</pre>	
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 Config Interface VLAN Config Interface Loopback	
History	3.4.2008	
Role	admin	
Example	<pre>switch (config interface ethernet 1/2) # vrf forwarding my-vrf switch (config interface ethernet 1/2) #</pre>	
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	<pre>switch (config) # show routing-context vrf VRF active context: my-vrf switch (config) #</pre>
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.	

6.1.5.2 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 a router port.

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 in a VLAN.
The no form of the command disables L2 802.1Q encapsulation of traffic on a specified router port in a VLAN.

Syntax Description	vlan-id	Enables L2 802.1Q encapsulation of traffic on a router port 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		

6.1.5.3 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	ip routing vlan <vlan-id> switchport mode switchport access show interfaces vlan	
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	<pre>switch (config interface vlan 10) # ip address 10.10.10.10 /24 switch (config interface vlan 10) #</pre>	
Related Commands	<pre>interface vlan show interfaces vlan</pre>	
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	<pre>switch (config interface vlan 10) # ip address dhcp switch (config interface vlan 10) #</pre>
Related Commands	<pre>interface vlan show interfaces vlan</pre>
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	<pre>switch (config interface vlan 10) # counters switch (config interface vlan 10) #</pre>
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	<pre>switch (config interface vlan 10) # description my-ip-interface switch (config interface vlan 10) #</pre>	
Related Commands	<pre>interface vlan show interfaces vlan</pre>	
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	<pre>switch (config interface vlan 10)# mtu 9216 switch (config interface vlan 10 #</pre>	
Related Commands	<pre>interface vlan show interfaces vlan</pre>	
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	<pre>switch (config interface vlan 10) # clear counters switch (config interface vlan 10) #</pre>
Related Commands	<pre>interface vlan counters</pre>
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.					

6.1.5.4 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	<pre>switch (config) # interface loopback 10 switch (config interface loopback 10) #</pre>	
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	<pre>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) #</pre>	
Related Commands		
Note		

6.1.5.5 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 • destination-ip-port • source-destination-ip-port • traffic-class • all
Default	all
Configuration Mode	Config
History	3.2.0230
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	

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	

6.1.5.6 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: 60-28800.
Default	1500 seconds	
Configuration Mode	Config Interface VLAN	
History	3.2.0230	
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	<pre>ip arp show ip arp</pre>	
Note	This value is used as the ARP timeout whenever a new IP interface is created.	

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	<pre>switch (config) # clear ip arp vrf my-vrf switch (config) #</pre>	
Related Commands	<pre>ip arp show ip arp</pre>	
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.	

6.1.5.7 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
	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 Added VRF parameter
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 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) #</pre>
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.	

6.1.5.8 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	<pre>switch (config) # qos map dscp-to-pcp preserve-pcp switch (config) #</pre>
Related Commands	
Note	• This commands applies the configuration for all router interfaces • As part of its function, the router performs DSCP to PCP bits mapping (fixed mapping). By activating the command the router preserves the PCP bits from one subnet to another subnet (PCP bits are copied).

6.2 IPv6



To activate this feature please contact your Mellanox support representative.

IP version 6 (IPv6) is a routing protocol which succeeds IPv4. With the expansion of the Internet and data bases IPv6 addresses consist of 128 bits whose purpose is to allow networks to include a significantly higher number of nodes by increasing the pool of available unique IP addresses. IPv6 packets alleviate overhead and allow for future customizability.

Textual representations of IPv6 addresses consist of 128 bits made up from eight 16-bit hexadecimal numbers separated by colons. IPv6 addresses may be abbreviated as follows:

- You may omit leading zeros in each 16-bit sequence
- You may replace an entire sequence with a double colon if it equals zero

For example, these addresses represent the same IPv6 address:

- af23:0000:0000:0000:1284:037d:35ce:2401
- af23:0:0:0:1284:37d:35ce:2401
- af23::1284:37d:35ce:2401

IPv6 addresses typically denote a 64-bit network prefix and a 64-bit host address.



Only static IPv6 and ECMP are supported.



The number of static IPv6 addresses supported is 64.

6.2.1 Neighbor Discovery Protocol

Neighbor Discovery (ND) decides relationships between neighbors and replaces ARP, ICMP, and ICMP redirect in IPv4.


Five kinds of ICMPv6 packets are defined by ND:

- Neighbor advertisement
- Router advertisement
- Neighbor solicitation
- Router solicitation
- Redirect

ND checks whether a neighboring node's address has changed, whether the neighbor is still reachable, and also resolves the address of the neighbor which a packet is being forwarded to. ND is also useful for network nodes for discovering other nodes and performing basic link-layer configuration.

6.2.2 Configuring IPv6

Figure 28: IPv6 Network



➤ To configure Router1:

Step 1. Enable IP routing. Run:

```
switch (config)# ip routing
```

Step 2. Enable forwarding IPv6 unicast packets. Run:

```
switch (config)# ipv6 routing
```

Step 3. Configure the VLAN interfaces. Run:

```
switch (config)# interface vlan 10
switch (config interface vlan 10) # exit
switch (config)# interface vlan 30
switch (config interface vlan 30) # exit
switch (config)# interface vlan 50
switch (config interface vlan 50) # exit
```

Step 4. Enable IPv6 on the VLAN interfaces. Run:

```
switch (config)# interface vlan 10 ipv6 enable
switch (config)# interface vlan 30 ipv6 enable
switch (config)# interface vlan 50 ipv6 enable
```

Step 5. Configure IPv6 addresses for each one of the VLAN interfaces. Run:

```
switch (config)# interface vlan 10 ipv6 address 2101:db01::1 /64
switch (config)# interface vlan 30 ipv6 address 2103:db01::2 /64
switch (config)# interface vlan 50 ipv6 address 2105:db01::1 /64
```

Step 6. Configure IPv6 unicast. Run:

```
switch (config)# ipv6 route 2002:db01:: /64 2101:db01::2
```

Step 7. Configure IPv6 unicast. Run:

```
switch (config)# ipv6 route 2002:db01:: /64 2105:db01::2
```

➤ To configure Router2:

Step 1. Disable prefix mode on the CLI. Run:

```
switch (config)# no cli default prefix-mode enable
```

Step 2. Enable the VLANs on the system. Run:

```
switch (config)# vlan 10
switch (config vlan 10) # exit
switch (config)# vlan 20
switch (config vlan 20) # exit
switch (config)# vlan 50
switch (config vlan 50) # exit
```

Step 3. Configure the switch ports to accept the VLANs of which they are part only. Run:

```
switch (config)# interface ethernet 1/1 switchport access vlan 10 // port2
switch (config)# interface ethernet 1/2 switchport access vlan 50 // port8
switch (config)# interface ethernet 1/36 switchport access vlan 20 // port5
```

Step 4. Disable spanning tree. Run:

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

Step 5. Enable IP routing. Run:

```
switch (config)# ip routing
```

Step 6. Enable forwarding IPv6 unicast packets. Run:

```
switch (config)# ipv6 routing
```

Step 7. Configure the VLAN interfaces. Run:

```
switch (config)# interface vlan 10
switch (config interface vlan 10) # exit
switch (config)# interface vlan 20
switch (config interface vlan 20) # exit
switch (config)# interface vlan 50
switch (config interface vlan 50) # exit
```

Step 8. Enable IPv6 on the VLAN interfaces. Run:

```
switch (config)# interface vlan 10 ipv6 enable
switch (config)# interface vlan 20 ipv6 enable
switch (config)# interface vlan 50 ipv6 enable
```

Step 9. Configure IPv6 addresses for each one of the VLAN interfaces. Run:

```
switch (config)# interface vlan 10 ipv6 address 2101:db01::2 /64
switch (config)# interface vlan 20 ipv6 address 2102:db01::1 /64
switch (config)# interface vlan 50 ipv6 address 2105:db01::2 /64
```

Step 10. Configure IPv6 unicast. Run:

```
switch (config)# ipv6 route 2103:db01:: /64 2101:db01::1
```

Step 11. Configure IPv6 unicast. Run:

```
switch (config)# ipv6 route 2103:db01:: /64 2105:db01::1
```

➤ ***Ping neighbor to verify IPv6 configuration:***

```
switch (config)# ping6 2101:db01::2
PING 2101:db01::2(2101:db01::2) 56 data bytes
64 bytes from 2101:db01::2: icmp_seq=1 ttl=64 time=0.371 ms
64 bytes from 2101:db01::2: icmp_seq=2 ttl=64 time=0.620 ms
64 bytes from 2101:db01::2: icmp_seq=3 ttl=64 time=0.192 ms
64 bytes from 2101:db01::2: icmp_seq=4 ttl=64 time=0.277 ms
64 bytes from 2101:db01::2: icmp_seq=5 ttl=64 time=0.231 ms
```

6.2.3 Commands

ipv6 enable

ipv6 enable
no ipv6 enable

Assigns automatic local IPv6 address to the interface.
 The no form of the command deassigns that automatic local address and disables IPv6 if no static IPv6 address has been assigned to the interface.

Syntax Description	N/A
Default	Unassigned
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port
History	3.4.1100
Role	admin
Example	switch (config vlan 10) # ipv6 enable
Related Commands	
Note	Assigning an IPv6 address to an interface enables IPv6 processing on the interface also

ipv6 address

ipv6 address <ipv6-address> [/<length>
no ipv6 address <ipv6-address> [/<length>]

Enables IPv6 processing and assigns an IPv6 address to the interface.
 The no form of the command removes the specified IPv6 address.

Syntax Description	ipv6-address	IPv6 address. Format: a:b:c:d:e:f:g:h.
	length	Mask length for the associated address space. Range: 1-128.
Default	N/A	
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port	
History	3.4.1100	
Role	admin	
Example	switch (config vlan 10) # ipv6 address fe80:ac10::fa01:0202 /120 switch (config vlan 10) # ipv6 address fe80:ac10::fa01:0202/120	
Related Commands		
Note	- An interface can have up to 16 IPv6 address assignments - If the no command does not include a specific address, all address assignments are removed from the interface - The mask length may be configured without a space (i.e. <ipv6-address>/<length>)	

ipv6 nd managed-config-flag

ipv6 nd managed-config-flag
no ipv6 nd managed-config-flag

Sets the managed address configuration flag in IPv6 router advertisements.
 The no form of the command restores the default setting.

Syntax Description	N/A
Default	Managed address configuration flag is not set
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port
History	3.4.1100
Role	admin
Example	switch (config vlan 10) # ipv6 nd managed-config-flag
Related Commands	
Note	

ipv6 nd ns-interval

ipv6 nd ns-interval <period>

no ipv6 nd ns-interval

Configures the interval between IPv6 neighbor solicitation (NS) transmissions.
The no form of the command restores the default value.

Syntax Description	period	In milliseconds. Range: 1000-4294967295.
Default	1000 milliseconds	
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port	
History	3.4.1100	
Role	admin	
Example	switch (config vlan 10) # ipv6 nd ns-interval 1500	
Related Commands		
Note		

ipv6 nd other-config-flag

ipv6 nd other-config-flag
no ipv6 nd other-config-flag

Indicates that other configuration information is available via DHCPv6.
 The no form of the command removes the other configuration flag.

Syntax Description	N/A
Default	Not set
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port
History	3.4.1100
Role	admin
Example	switch (config vlan 10) # ipv6 nd other-config-flag
Related Commands	
Note	

ipv6 nd prefix

ipv6 nd prefix <ipv6-address> /<length> [no-advertise] [no-autoconfig] [no-onlink] [valid-time {<time> | infinite}] [preferred-time {<time> | infinite}]
no ipv6 nd prefix <prefix>

Configures inclusion for router advertisements (RAs) for neighbor.
 The no form of the command removes the corresponding IPv6 nd prefix.

Syntax Description	ipv6-address	IPv6 address. Format: a:b:c:d:e:f:g:h.
	length	Prefix length for the associated address space. Range: 1-128.
	no-advertise	Prevents advertising of the specified prefix.
	valid-time	Time in seconds. Range: 0-4294967295. The value “infinite” is the same as the maximum value possible.
	preferred-time	Time in seconds. Range: 0-4294967295. The value “infinite” is the same as the maximum value possible.
	no-autoconfig	Indicates that this prefix can be used for stateless address configuration
	no-onlink	Indicates that this prefix can be used for on-link determination
Default	valid-time: 2592000 seconds preferred-time: 604800 seconds no-autoconfig: autoconfig enabled no-onlink: Set	
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port	
History	3.4.1100	
Role	admin	
Example	switch (config vlan 10) # ipv6 nd prefix fe80:ac10::fa01:0202 /120	
Related Commands		
Note	Valid time must be larger than preferred time	

ipv6 nd ra dns-servers lifetime

ipv6 nd ra dns-servers lifetime {<time> | infinite}

no ipv6 nd ra dns-servers lifetime

Sets the default value for the lifetime of any recursive DNS server (RDNSS) configured on the interface.

The no form of the command removes the lifetime value.

Syntax Description	time	Possible values: 0 – RDNSS configured on the command mode interface without a custom lifetime value must not be used 1-4294967295 – in seconds
	infinite	Equivalent to 4294967295.
Default	If no lifetime period is configured on the interface, the default value is 1.5 times the maximum RA interval set by the command “ipv6 nd ra interval”.	
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port	
History	3.4.1100	
Role	admin	
Example	switch (config vlan 10) # ipv6 nd ra dns-servers lifetime infinite	
Related Commands		
Note	- A lifetime value set for an individual RDNSS overrides this value. - The lifetime value is the maximum amount of time after a route advertisement packet is sent that the RDNSS referenced in the packet may be used for name resolution.	

ipv6 nd ra dns-server

ipv6 nd ra dns-server <ip-address> [lifetime [<time> | infinite]]
no ipv6 nd ra dns-server [<ip-address>]

Configures the IPv6 address of a preferred recursive DNS server (RDNSS) to include in the neighbor-discovery router advertisements (RAs).

The no form of the command removes the RDNSS from the configuration.

Syntax Description	ip-address	IP address of RDNSS
	lifetime	Maximum lifetime value for the specified RDNSS entry. Possible values: • 0 – RDNSS address must no longer be used • 1-4294967295 in seconds
	infinite	Equivalent to 4294967295 seconds.
Default	If no lifetime period is configured on the interface, the default value is 1.5 times the maximum RA interval set by the command “ipv6 nd ra interval”.	
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port	
History	3.4.1100	
Role	admin	
Example	<pre>switch (config vlan 10) # ipv6 nd ra dns-server fe80:ac10::fa01:0202 lifetime infinite</pre>	
Related Commands		
Note	• Including RDNSS information in RAs provides DNS server configuration for connected IPv6 hosts without requiring DHCPv6 • Multiple servers can be configured on the interface by using the command repeatedly • A lifetime value for the RDNSS can optionally be specified with this command, and overrides any default value configured for the interface using the ipv6 nd ra dns-servers lifetime command • Lifetime must be configured according to the following: $\text{MaxRtrAdvInterval} \leq \text{lifetime} \leq 2 * \text{MaxRtrAdvInterval}$; where MaxRtrAdvInterval is the maximum RA interval	

ipv6 nd ra dns-suffix

ipv6 nd ra dns-suffix <domain-name> [lifetime {<time> | infinite}]
no ipv6 nd ra dns-suffix [<domain-name>]

Creates a DNS search list (DNSSL) to include in the neighbor-discovery router advertisements (RAs)¹.

The no form of the command resets the value of this parameter to its default.

Syntax Description	domain-name	Domain suffix for IPv6 hosts to append to short unqualified domain names for DNS queries. The suffix must contain only alphanumeric characters, “.” (periods), “-” (hyphens), and must begin and end with an alphanumeric character.
	lifetime	Sets the maximum time, in seconds (relative to the time the packet is sent), over which this DNSSL domain name MAY be used for name resolution. Lifetime must be bounded as follows: MaxRtrAdvInterval <= Lifetime <= 2*MaxRtrAdvInterval
	time	Possible values: • 0 – DNSSL must not be used for name resolution • 1-4294967295 – in seconds
	infinite	A value of all one bits (0xffffffff) equivalent to 4294967295.
Default	The value configured in the command “ipv6 nd ra dns-server”.	
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port	
History	3.4.1100	
Role	admin	
Example	<pre>switch (config vlan 10) # ipv6 nd ra dns-suffix example lifetime infinite</pre>	
Related Commands		
Note	• The DNSSL contains the domain names of DNS suffixes for IPv6 hosts to append to short, unqualified domain names for DNS queries • Multiple DNS domain names can be added to the DNSSL by reusing the command • A lifetime value for the DNSSL can optionally be specified with this command which overrides any default value configured for the interface using the command “ipv6 nd ra dns-suffixes lifetime”	

1. As defined in RFC 6106.

ipv6 nd ra dns-suffixes lifetime

ipv6 nd ra dns-suffixes lifetime {<time> | infinite}

no ipv6 nd ra dns-suffixes lifetime

Creates a DNS search list (DNSSL) to include in the neighbor-discovery router advertisements (RAs)¹.

The no form of the command resets the value of this parameter to its default.

Syntax Description	domain-name	Domain suffix for IPv6 hosts to append to short unqualified domain names for DNS queries. The suffix must contain only alphanumeric characters, "." (periods), "-" (hyphens), and must begin and end with an alphanumeric character.
	time	Possible values: 0 – DNSSL must not be used for name resolution if no custom value is configured 1-4294967295 – in seconds
	infinite	Equivalent to 4294967295.
Default	1.5 times the RA interval configured on the interface	
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port	
History	3.4.1100	
Role	admin	
Example	<pre>switch (config vlan 10) # ipv6 nd ra dns-suffix example lifetime infinite</pre>	
Related Commands		
Note	- The DNSSL contains the domain names of DNS suffixes for IPv6 hosts to append to short, unqualified domain names for DNS queries - Multiple DNS domain names can be added to the DNSSL by reusing the command	

1. As defined in RFC 6106.

ipv6 nd ra hop-limit

ipv6 nd ra hop-limit <limit>

no ipv6 nd ra hop-limit

Sets a suggested hop-limit value to be included in route advertisement (RA) packets. The no form of the command resets the parameter to its default value.

Syntax Description	limit	The hop-limit value to be included by attached hosts in outgoing packets. • 0 – unspecified (by this router) • 1-255 – number of hops
Default	Limit value is 64	
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port	
History	3.4.1100	
Role	admin	
Example	switch (config vlan 10) # ipv6 nd ra hop-limit 70	
Related Commands		
Note		

ipv6 nd ra interval max-period

ipv6 nd ra interval max-period <time> [min-period <time>]
no ipv6 nd ra interval

Configures the interval between IPv6 router advertisement (RA) transmissions.
 The no form of the command resets the parameter to its default value.

Syntax Description	time	Maximum interval between successive IPv6 router advertisement transmissions. Range: 4-1800 seconds.
	min-period	minimum interval between successive IPv6 router advertisement transmissions. - No parameter: Default is used 4-1800: Valid range when scale is set seconds.
Default	max-period: 600 seconds min-period: See Note	
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port	
History	3.4.1100	
Role	admin	
Example	switch (config vlan 10) # ipv6 nd ra interval max-period 600	
Related Commands		
Note	- The min-period must be $0.33 * \text{<max-period>}$ if <max-period> is ≥ 9 seconds; otherwise, the default is MaxRtrAdvInterval - The parameter min-period must be no less than 3 seconds and no greater than $0.75 * \text{max-period}$	

ipv6 nd ra lifetime

ipv6 nd ra lifetime <time>

no ipv6 nd ra lifetime

Configures the value that the switch places in the router lifetime field of IPv6 router advertisements (RAs).

The no form of the command resets the parameter to its default value.

Syntax Description	time	The router lifetime specifies the period that the router can be considered as a default router by RA recipients in seconds. 0 – the router should not be considered a default router on this interface 1-9000 – lifetime period advertised in RAs should not be less than the max router advertisement interval
Default	3*<max router advertisement interval>	
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port	
History	3.4.1100	
Role	admin	
Example	switch (config vlan 10) # ipv6 nd ra lifetime 300	
Related Commands		
Note		

ipv6 nd ra mtu suppress

ipv6 nd ra mtu suppress
no ipv6 nd ra mtu suppress

Suppresses the router advertisement (RA) MTU option to ensure that all nodes on a link use the same MTU value.
 The no form of the command restores the MTU option to enabled.

Syntax Description	N/A
Default	Suppressed
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port
History	3.4.1100
Role	admin
Example	switch (config vlan 10) # ipv6 nd ra mtu suppress
Related Commands	
Note	If not suppressed, MTU of the interface is advertised.

ipv6 nd ra suppress

ipv6 nd ra suppress [all]
no ipv6 nd ra suppress

Suppresses periodic IPv6 router advertisement (RA) transmissions.
 The no form of the command restores the transmission of RAs.

Syntax Description	all	Configures the switch to suppress all RAs, including those responding to a router solicitation.
Default	Only unsolicited RAs transmitted periodically are suppressed	
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port	
History	3.4.1100	
Role	admin	
Example	switch (config vlan 10) # ipv6 nd ra suppress all	
Related Commands		
Note		

ipv6 nd reachable-time

ipv6 nd reachable-time <time>

no ipv6 nd reachable-time

Sets the time period the switch includes in the reachable time field of out-going advertisements (RAs).

The no form of the command resets the parameter to its default value.

Syntax Description	time	In milliseconds; the reachable time defines the period that a node assumes a neighbor is reachable after having received a reachability confirmation. Range: 0-3600000 where “0” means unspecified by this router.
Default	0 (unspecified)	
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port	
History	3.4.1100	
Role	admin	
Example	switch (config vlan 10) # ipv6 nd reachable-time 30000	
Related Commands		
Note	• RAs that advertise zero seconds indicate that the router does not specify a reachable time • The default value to use for calculating neighbor reachability time is 30 seconds	

ipv6 nd router-preference

ipv6 nd router-preference {high | medium | low}

no ipv6 nd router-preference

Sets the value the switch enters in the default router preference (DRP) field of router advertisements (RAs) it sends.

The no form of the command resets the parameter to its default value.

Syntax Description	N/A
Default	Medium
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port
History	3.4.1100
Role	admin
Example	<pre>switch (config vlan 10) # ipv6 nd router-preference high</pre>
Related Commands	
Note	- IPv6 hosts maintain a default router list from which to select a router for traffic to offlink destinations. The router's address is then saved in the destination cache. The neighbor discovery protocol (NDP) prefers routers that are reachable or probably reachable over routers whose reachability is unknown or suspect. For reachable or probably reachable routers, NDP can either select the same router every time or cycle through the router list. DRP values specify a host's preferred router. - If router lifetime is zero, preference value must be medium

ipv6 nd retrans-timer

ipv6 nd retrans-timer <time>

no ipv6 nd retrans-timer

Advertises the time between neighbor solicitation (NS) messages in ICMPv6 router advertisement messages.

The no form of the command resets the parameter to its default value.

Syntax Description	time	In milliseconds; the time between retransmitted neighbor solicitation messages. Possible values: • 0 – unspecified • Range – 1000-4294967295
Default	0 (unspecified)	
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port	
History	3.4.1100	
Role	admin	
Example	switch (config vlan 10) # ipv6 nd retrans-timer	
Related Commands		
Note		

ipv6 nd dad attempts

ipv6 nd dad attempts <number>

no ipv6 nd dad attempts

Sets the number of consecutive neighbor solicitation messages sent for duplicate address detection (DAD) validation.

The no form of the command resets the value to its default.

Syntax Description	number	Number of attempts: • 0 – DAD is not performed • Valid range: 1-1000
Default	1	
Configuration Mode	Config Interface VLAN Config Interface Loopback Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port	
History	3.4.1100	
Role	admin	
Example	switch (config vlan 10) # ipv6 nd dad attempts 10	
Related Commands		
Note		

ipv6 neighbor

```

ipv6 neighbor <ipv6-addr> {ethernet <port> | port-channel <port-channel> |
vlan <vlan-id>} <mac_addr>
no ipv6 neighbor <ipv6-addr> [{ethernet <port> | port-channel <port-channel> |
vlan <vlan-id>}]

```

Creates an IPv6 neighbor discovery cache static entry.
 The no form of the command removes the specified static entry from the IPv6 neighbor discovery cache.

Syntax Description	ipv6-addr	IPv6 address
	ethernet <port>	Ethernet port. Format <slot>/<port>.
	vlan <vlan-id>	VLAN ID
Default	N/A	
Configuration Mode	Config	
History	3.4.1100	
Role	admin	
Example	switch (config vlan 10) # ipv6 neighbor 2001:db01::1 vlan 10 4:4:4:4:4:4	
Related Commands		
Note	This command do not affect any dynamic entries in the cache.	

clear ipv6 neighbors

clear ipv6 neighbors [ethernet <port> | port-channel <port-channel> | vlan <vlan-id>] [<ipv6-addr>]

Removes the specified dynamic IPv6 neighbor discovery cache entries.

Syntax Description	ethernet	Ethernet port. Format: <slot>/<port>.
	vlan	VLAN interface
	ipv6-addr	IPv6 address
Default	N/A	
Configuration Mode	Config	
History	3.4.1100	
Role	admin	
Example	switch (config) # clear ipv6 neighbors ethernet 1/4	
Related Commands		
Note	• Commands that do not specify an IPv6 address remove all dynamic entries for the listed interface • Commands that do not specify an interface remove all dynamic entries • See the command “clear ipv6 neighbors” on page 139	

ipv6 route

ipv6 route <ipv6-address> /<length> <next-hop> [<distance>]
no ipv6 route <ipv6-address> /<length> [next-hop]

Creates an IPv6 static route.
 The no form of the command deletes static routes.

Syntax Description	ipv6-address	IPv6 address. Format: a:b:c:d:e:f:g:h.
	length	Prefix length for the associated address space. Range: 1-128.
	next-hop	IPv6 address of the next-hop
	distance	Administrative distance assigned to route. Options include: - No parameter – route is assigned a default administrative distance of 1 1-255 – the administrative distance assigned to route
Default	No distance parameter indicated: Administrative distance of 1	
Configuration Mode	Config	
History	3.4.1100	
Role	admin	
Example	<pre>switch (config) # ipv6 route 3003:db01:: /64 2001:db01::1 switch (config) #</pre>	
Related Commands		
Note	- Static routes have a default administrative distance of 1 - Assigning a higher administrative distance to a static route configures it to be overridden by dynamic routing data. - Multiple routes which are configured to the same destination with the same administrative distance comprise an Equal Cost Multi-Path (ECMP) route - A no command not including a source deletes all statements to the destination	

ipv6 routing

ipv6 routing
no ipv6 routing

Enables forwarding IPv6 unicast packets.
 The no form of the command disables IPv6 unicast routing.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config
History	3.4.1100
Role	admin
Example	<code>switch (config) # ipv6 routing</code>
Related Commands	
Note	When routing is enabled, the switch attempts to deliver inbound packets to destination addresses by forwarding them to interfaces or next hop addresses specified by the IPv6 routing table

show ipv6 interface

show ipv6 interface [{**ethernet** <port> | **port-channel** <port-channel> | **vlan** <vlan-id>}] [**brief**]

Displays the status of specified routed interfaces that are configured for IPv6.

Syntax Description	ethernet <port>	Displays output pertaining to the specified Ethernet interface
	port-channel <port-channel>	Displays output pertaining to the specified LAG interface
	vlan <vlan-id>	Displays output pertaining to the specified VLAN interface
	brief	Shows basic IPv6 information regarding all IPv6 interfaces
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.4.1100	
Role	admin	
Example	<pre>switch (config) # show ipv6 interface Vlan10 is Enabled , line protocol is UP IPv6 : Enabled Link-local address : fe80::f652:14ff:fe2d:9808 Global Unicast Addresses : 2001:db01::2 /64 Joined Group Addresses : ff02::1 ff02::2 ff02::1:ff2d:9808 MTU : 1500 bytes ICMP error messages limited to every milliseconds : 100 ICMP redirects : enabled ND DAD : enabled Number of DAD attempts : 1 ND reachable time (milliseconds) : 30000 ND advertised retransmit interval (milliseconds) : 0 ND router advertisements maximum interval (seconds) : 600 ND router advertisements minimum interval (seconds) : 198 ND router advertisements managed configuration flag : unset ND router advertisements other configuration flag : unset ND solicited router advertisement : suppressed ND router advertisements lifetime (seconds) : 1800 ND advertised default router preference : medium ND router advertisements hop-limit : 64 switch (config) #</pre>	
Related Commands		
Note		

show ipv6 neighbors

show ipv6 neighbors [{**ethernet** <port> | **port-channel** <port-channel> | **vlan** <vlan-id>} | <ipv6-addr> | **summary**]

Displays IPv6 neighbor discovery (ND) cache information.

Syntax Description	ethernet <port>	Shows output pertaining to the specified Ethernet interface.
	vlan <vlan-id>	Shows output pertaining to the specified VLAN interface.
	ipv6-addr	IPv6 address of individual neighbor
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.4.1100	
Role	admin	
Example	<pre>switch (config) # show ipv6 route IPv6 Address MAC Address State Interf ----- 2001:db01::1 f4:52:14:2d:98:88 Reachable vlan10 switch (config) #</pre>	
Related Commands		
Note		

show ipv6 route

show ipv6 route [<ipv6-addr> <prefix>] [connected | static | summary]

Displays IPv6 neighbor discovery (ND) cache information.

Syntax Description	ipv6-addr	Filters routes by IPv6 address or prefix
	longer-prefixes	Displays output for longer prefix entries
	connected	Displays entries for routes to networks directly connected to the switch
	static	Displays entries added through CLI commands
	summary	Displays the current contents of the IPv6 routing table in summary format
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.4.1100	
Role	admin	
Example	<pre>switch (config) # show ipv6 route Destination Mask Gateway Interface Source Distance/Metric fe80:: 64 :: mgmt0 Connected 256/1 fe80:: 64 :: mgmt1 Connected 256/1 2001:db01:: 64 :: vlan10 Connected 1/1 3003:db01:: 64 2001:db01::1 vlan10 Static 1/20 switch (config) #</pre>	
Related Commands		
Note		

6.3 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](#).

6.3.1 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

6.3.2 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.

6.3.3 Configuring OSPF

Figure 29: OSPF Basic Topology



Precondition steps:



The following configuration example refers to Router 2 in [Figure 29](#). 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 31](#).

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
```

6.3.4 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) #
```

6.3.5 Commands

6.3.5.1 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
no router ospf

Enters router OSPF configuration mode, and creates default OSPF instance if not exist.

The no form of the command deletes the OSPF instance.

Syntax Description	N/A
Default	No router OSPF is created.
Configuration Mode	Config
History	3.3.3500
Role	admin
Example	switch (config)# router ospf switch (config router ospf) #
Related Commands	N/A
Note	Only one OSPF instance is supported.

6.3.5.2 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	<code>switch (config router ospf)# shutdown</code>
Related Commands	N/A
Note	

auto-cost reference-bandwidth

auto-cost reference-bandwidth <ref-bw>

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	value	Range: 1-4294
Default	Gb/s	
Configuration Mode	Config OSPF Router	
History	3.3.3500	
Role	admin	
Example	switch (config router ospf)# auto-cost reference-bandwidth	
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 routes 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	Netowrk 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 advertized into the stub area.
Default	Summary route will be advertized.	
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 advertized 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.	

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.	

6.3.5.3 Interface

ip ospf cost

ip ospf cost <cost>
no ip ospf cost <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 Config Interface Port Channel configured as a router port	
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 Config Interface Port Channel configured as a router port	
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 Config Interface Port Channel configured as a router port	
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 Config Interface Port Channel configured as a router port	
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 Config Interface Port Channel configured as a router port	
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 Config Interface Port Channel configured as a router port	
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 Config Interface Port Channel configured as a router port
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 Config Interface Port Channel configured as a router port	
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 Config Interface Port Channel configured as a router port
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 Config Interface Port Channel configured as a router port	
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 Config Interface Port Channel configured as a router port	
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 Config Interface Port Channel configured as a router port	
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 Config Interface Port Channel configured as a router port Config Interface Loopback	
History	3.3.3500	
Role	admin	
Example	switch (config interface vlan 10)# ip ospf area 0	
Related Commands	N/A	
Note		

6.3.5.4 Show

show ip ospf

show ip ospf

Displays general OSPF configuration and status.

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 Routing Process 201 with ID 192.0.2.1 VRF default Admin Status is Enabled Stateful High Availability enabled Graceful-restart is configured</pre>
Related Commands	N/A
Note	

show ip ospf border-routers

show ip ospf border-routers

Displays routing table entries to an Area Border Routers.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.3500
Role	admin
Example	<pre>switch# show ip ospf border-routers OSPF Process ID p1, vrf default Internal Routing Table Codes: i - Intra-area route, I - Inter-area route i 40.40.40.40 [10], ABR, Area 0.0.0.0, SPF 71 via 192.0.2.1, Ethernet2/1 i 60.60.60.60 [20], ABR, Area 0.0.0.0, SPF 71 via 192.0.2.1, Ethernet2/1 i 40.40.40.40 [10], ABR, Area 0.0.0.1, SPF 71 via 192.0.2.1, Ethernet2/2 i 60.60.60.60 [20], ABR, Area 0.0.0.1, SPF 71 via 192.0.2.1, Ethernet2/2</pre>
Related Commands	N/A
Note	

show ip ospf database

show ip ospf database [summary] [<area-id> [<link-state-id>]] [adv-router <ip-address> | self-originated]

Displays the OSPF database.

Syntax Description	adv-router <ip-address>	Filters per advertize 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.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.3500	
Role	admin	
Example	<pre>Router# show ip ospf database OSPF Router with ID (50.50.50.50) (Process ID p1) Router Link States (Area 0) Link ID ADV Router Age Seq# Checksum Link Count 40.40.40.40 40.40.40.40 930 0x80000004 0x2ea1 3 50.50.50.50 50.50.50.50 935 0x80000002 0x8b52 1 60.60.60.60 60.60.60.60 943 0x800003c5 0x9854 2 Network Link States (Area 0) Link ID ADV Router Age Seq# Checksum 209.165.201.3 60.60.60.60 944 0x80000001 0x7179 192.0.2.1 50.50.50.50 935 0x80000001 0x516a Summary Network Link States (Area 0) Link ID ADV Router Age Seq# Checksum 209.165.201.1 40.40.40.40 929 0x80000001 0x2498 209.165.201.1 50.50.50.50 928 0x80000001 0x5b2f 209.165.201.1 60.60.60.60 1265 0x800003c3 0xf49b 192.0.2.0 40.40.40.40 943 0x80000001 0x53f3 192.0.2.0 50.50.50.50 935 0x80000001 0x26f8 192.0.2.0 60.60.60.60 930 0x80000001 0x7b51</pre>	
Related Commands	N/A	
Note		

show ip ospf interface

show ip ospf interface [vlan <vlan-id>] [brief]

Displays the OSPF related interface configuration.

Syntax Description	brief	Gives a brief summary of the output.
	vlan <vlan-id>	Displays OSPF interface configuration and status per VLAN interface.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.3500	
Role	admin	
Example	<pre>switch# show ip ospf interface ethernet 1/5 Ethernet1/5 is up, line protocol is down IP address 192.0.2.1, Process ID 201 VRF RemoteOfficeVRF, area 0.0.0.10 Enabled by interface configuration State DOWN, Network type BROADCAST, cost 4 Index 1, Transmit delay 1 sec, Router Priority 1 No designated router on this network No backup designated router on this network 0 Neighbors, flooding to 0, adjacent with 0 Timer intervals: Hello 10, Dead 40, Wait 40, Retransmit 5 No authentication Number of opaque link LSAs: 0, checksum sum 0 switch#</pre> This example shows how to display OSPF information in a brief format: <pre>switch# show ip ospf interface brief OSPF Process ID 201 VRF default Total number of interface: 1 Interface ID Area Cost State Neighbors Status VL1 2 0.0.0.0 65535 DOWN 0 down switch#</pre>	
Related Commands	N/A	
Note		

show ip ospf neighbors

show ip ospf neighbors [vlan <vlan-id>] [<neighbor-id>]

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.
Default	N/A	
Configuration Mode	Any Command Mode	
History	3.3.3500	
Role	admin	
Example	<pre>Router# show ip ospf neighbors 10.199.199.137 Neighbor 10.199.199.137, interface address 192.0.2.37 In the area 0.0.0.0 via interface Ethernet2/1 Neighbor priority is 1, State is FULL Options 2 Dead timer due in 0:00:32 Link State retransmission due in 0:00:04 Neighbor 10.199.199.137, interface address 209.165.201.189 In the area 0.0.0.0 via interface Ethernet4/3 Neighbor priority is 5, State is FULL Options 2 Dead timer due in 0:00:32 Link State retransmission due in 0:00:03</pre> This example shows how to display the neighbors that match the neighbor ID on an interface: <pre>Router# show ip ospf neighbors ethernet 2/1 10.199.199.137 Neighbor 10.199.199.137, interface address 192.0.2.37 In the area 0.0.0.0 via interface Ethernet2/1 Neighbor priority is 1, State is FULL Options 2 Dead timer due in 0:00:37 Link State retransmission due in 0:00:04 This example shows how to display detailed information about OSPF neighbors:</pre> <pre>Router# show ip ospf neighbors detail Neighbor 192.168.5.2, interface address 10.225.200.28 In the area 0 via interface GigabitEthernet1/0/0 Neighbor priority is 1, State is FULL, 6 state changes DR is 10.225.200.28 BDR is 10.225.200.30 Options is 0x42 LLS Options is 0x1 (LR), last OOB-Resync 00:03:08 ago Dead timer due in 00:00:36 Neighbor is up for 00:09:46 Index 1/1, retransmission queue length 0, number of retransmission 1 First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0) Last retransmission scan length is 1, maximum is 1 Last retransmission scan time is 0 msec, maximum is 0 msec</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</pre> Display of Summary addresses for External Routes and area ranges for the summary LSAs <pre>OSPF Process default OSPF External Summary Address and area-range Configuration Information -----</pre> <table><tr><th>Network</th><th>Mask</th><th>Area</th><th>Advertise</th><th>LSA type</th><th>Metric</th><th>Tag</th></tr><tr><td>1.1.1.1</td><td>255.255.255.0</td><td>NA</td><td>Advertise</td><td>Type5</td><td>10</td><td>0</td></tr><tr><td>2.2.2.0</td><td>255.255.255.0</td><td>10.10.10.10</td><td>Not Advertise</td><td>Type3</td><td>10</td><td>0</td></tr></table>	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
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																
Related Commands	N/A																					
Note																						

6.4 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.


6.4.1 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.

6.4.2 Configuring BGP

Figure 30: 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
```

6.4.3 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.

6.4.4 Commands

6.4.4.1 Config

protocol bgp

protocol bgp
no protocol bgp

Enables BGP feature, 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 | out]]

Clears BGP learned routes from the BGP table and resets the connection to the neighbor.

Syntax Description	ip-address	A BGP peer IP address. Only the specified neighbor is reset.
	all	All BGP peers. All BGP neighbors are reset.
	soft	Clears BGP learned routes from the BGP table without resetting the connection to the neighbor.
	in	Inbound routes are reset.
	out	Outbound routes are reset.
Default	N/A	
Configuration Mode	Config	
History	3.3.5006	First release
	3.3.5200	Updated description
Role	admin	
Example	<pre>switch (config)# clear ip bgp all switch (config)#</pre>	
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	<pre>switch (config)# router bgp 100 switch (config router bgp 100)#</pre>	
Related Commands	ip routing	
Note		

6.4.4.2 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	<code>switch (config router bgp 100)# no shutdown</code>
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	<pre>switch-e07c04 [standalone: master] (config router bgp 4) # aggregate- address 3.5.3.7 /32</pre>	
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
no bestpath as-path multipath-relax

Enables ECMP across AS paths.
 The no form of the command disables ECMP across AS paths.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config Router BGP
History	3.3.5006
	3.3.5200 Updated description and notes
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

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	
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.	

bgp redistribute-internal

bgp redistribute-internal
no bgp redistribute-internal

Enables iBGP redistribution into an interior gateway protocol (IGP).
 The no form of the command disables iBGP redistribution into an interior gateway protocol (IGP).

Syntax Description	ip-prefix	IP address
	length	Mask length (e.g. /24 or 255.255.255.254)
	peer-group-name	Peer group name
	remote-as <as-number>	Remote peer's number.
Default	Disabled	
Configuration Mode	Config Router BGP	
History	3.4.0000	
Role	admin	
Example	switch (config router bgp 100)# bgp redistribute-internal	
Related Commands		
Note		

cluster-id

cluster-id <ip-address>
no cluster-id <ip-address>

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”)
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
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
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

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.
Default	30 seconds	
Configuration Mode	Config Router BGP	
History	3.4.0000	First version
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: 1-10.
Default	N/A	
Configuration Mode	Config Router BGP	
History	3.4.0000	First version
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	<pre>switch (config router bgp 100)# neighbor 10.10.10.10 description The next door neighbor</pre>	
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 applies the system 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.
Default	ttl: 1	
Configuration Mode	Config Router BGP	
History	3.3.5006	First version
	3.3.5200	Updated default
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.
Default	100	
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 export-localpref 100</pre>	
Related Commands		
Note		

neighbor graceful-restart helper

neighbor {<ip-address> | <peer-group-name>} graceful-restart helper
no neighbor {<ip-address> | <peer-group-name>} graceful-restart helper

Enables BGP graceful restart helper mode for the specified BGP neighbor or peer group.

The no form of the command

Syntax Description	ip-address	IP address of the BGP-speaking neighbor
	peer-group-name	Peer group name
Default	Graceful restart is enabled	
Configuration Mode	Config Router BGP	
History	3.4.0000	First version
Role	admin	
Example	switch (config router bgp 100)# neighbor graceful-restart helper	
Related Commands		
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	

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.
Default	100	
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 import-localpref 100</pre>	
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
	no-prepend	local-as number is not prepended 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
Role	admin	
Example	<pre>switch-e07c04 [standalone: master] (config router bgp 4) # neighbor 100.100.100.100 local-as 123</pre>	
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.
	warning-only	Only generates a warning rather than disconnecting the neighbor
Default	12000	
Configuration Mode	Config Router BGP	
History	3.4.0000	First version
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
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
Default	no next-hop-peer	
Configuration Mode	Config Router BGP	
History	3.3.5006	
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
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
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	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 password 7 admin123</pre>	
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-text 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 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** <peer-group-name>

1. Assigns BGP neighbors to an existing peer group.
2. Creates a peer-group
3. Unassigns BGP neighbors to an existing peer group.
4. Removes a specified neighbor from 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
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	

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	<pre>switch (config router bgp 100)# neighbor 10.10.10.10 remote-as 200 switch (config router bgp 100)#</pre>	
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
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
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	ip routing router bgp <as-number>	
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
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.	

neighbor route-reflector-client

neighbor {<ip-address> | <peer-group-name>} route-reflector-client
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
Default	N/A	
Configuration Mode	Config Router BGP	
History	3.3.5006	First version
	3.3.5200	Updated notes and default
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
no neighbor {<ip-address> | <peer-group-name>} send-community

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
Default	Enabled	
Configuration Mode	Config Router BGP	
History	3.4.0000	First version
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
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
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.	

neighbor soft-reconfiguration inbound

neighbor {<ip-address> | <peer-group-name>} soft-reconfiguration inbound
no neighbor {<ip-address> | <peer-group-name>} soft-reconfiguration inbound

Disables BGP neighbor gracefully.

The no form of the command restores the system default behavior (retaining all routes from the specified neighbor or 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
Role	admin	
Example	switch (config router bgp 100)# neighbor 10.10.10.10 soft-reconfiguration inbound	
Related Commands	N/A	
Note	- This command also allows the switch to display all advertised routes when the command <code>show ip bgp neighbor advertised-routes</code> is issued. - The no form of the command configures the switch to discard information about routes received from the specified neighbor or group that fail the import policy.	

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.
	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.
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
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

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
Default	TCP sessions initiated	
Configuration Mode	Config Router BGP	
History	3.4.0000	First version
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 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.
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. For example, 192.0.2.4/24 is stored as 192.0.2.0/24.	

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.	

router-id

router-id <ip-address>

no router-id

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
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	
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.	

timers bgp

timers bgp <keep-alive> <hold>
no timers bgp

Configures the BGP keepalive and hold times.
 The no form of the command resets the parameters to their default settings.

Syntax Description	keep-alive	Frequency (in seconds) with which keepalive messages are sent to its peer. Range: 1-3600 seconds; 0 – no keep-alive messages are sent.
	hold	Interval (in seconds) after not receiving a keepalive message that a peer is declared dead. 3-7200 seconds; 0 – peer is held indefinitely regardless of keep-alive messages.
Default	Keepalive time: 60 secs Hold time: 180 secs	
Configuration Mode	Config Router BGP	
History	3.3.5006	First version
	3.3.5200	Updated syntax description, related commands and notes
Role	admin	
Example	<pre>switch (config router bgp 100)# timers bgp 61 181 switch (config router bgp 100)#</pre>	
Related Commands	ip routing neighbor timers router bgp <as-number> show ip bgp	
Note	• Timer settings apply to every peer connection. • The command “neighbor timers” configures the times on a specified peer connection. • Hold time should be three times longer than the keepalive setting.	

6.4.4.3 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 <ip> received switch-e07c04 [standalone: master] (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-e07c04 [standalone: master] (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.
	advertised	Displays routes advertised to the specified neighbor.
	received	Displays routes received and accepted from specified neighbor.
	both	Displays routes received from specified neighbor.
	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-e07c04 [standalone: master] (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-e07c04 [standalone: master] (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.
	received	Displays routes received and accepted from specified neighbor.
	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		
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 [tomere_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	

6.4.5 IP AS-Path Access-List

6.4.5.1 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	<pre>switch (config)# ip as-path access-list mylist permit switch (config)#</pre>	
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		

6.4.6 IP Community-List

6.4.6.1 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	<pre>switch (config)# ip community-list expanded mycommunity permit 1:[0-9]+</pre>	
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.	

6.5 Policy Rules

6.5.1 Route Map

Route maps define conditions for redistributing routes between routing protocols. A route map clause is identified by a name, filter type (permit or deny) and a sequence number. Clauses with the same name are components of a single route map; the sequence number determines the order in which the clauses are compared to a route.

6.5.1.1 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	<pre>switch (config) # route-map mymap permit 1200 switch (config route-map mymap permit 1200) #</pre>	
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	prefix-list-name
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	
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.	

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.	

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	<pre>switch (config route-map mymap permit 10)# match ip next-hop 10.10.10.10</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 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	
Role	admin	
Example	switch (config route-map mymap permit 10)# set community internet 1:3 additive	
Related Commands	N/A	
Note		

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	

6.5.2 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.

6.5.2.1 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		

6.6 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), 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.

6.6.1 Basic PIM-SM

PIM relies on the underlying topology gathering protocols that collect unicast routing information and build multicast routing information base (MRIB). The primary role of MRIB is to determine the next hop for PIM messages. MC data flows along with the reverse path of the PIM control.

MC tree construction contains three phases:

1. Construction of a shared distribution tree. This tree is built around a special designated router (DR) called the rendezvous point (RP).
2. Establishing a native forwarding path from MC sources to the RP
3. Building an optimized MC distribution tree from each MC source to all MC targets bypassing the RP

The first stage of the multicast tree establishment starts when the MC receiver expresses desire to start receiving MC data. It can happen as a result of using one of the L2 protocols like MLD or IGMP, or by static configuration. When such request is received by the last hop router (a designated router) this router starts to build a distribution path from the RP. It starts to send periodic “Join” messages to the nearest PIM neighbor router towards the RP. The next router continues to do the same. Eventually the process converges when Join messages reach RP or a router that has already created that distribution tree. Usually that tree is called a shared tree because it is created for any source for specific MC group G and is noted as (*,G).

At that stage, MC senders can start sending MC data. The DR next to the MC source extracts the packets from the data flow and tunnels them to the RP. The RP decapsulates the packets and distributes them to all MC receivers along with the share tree.

On the second stage the RP switches from tunneling of multicast packets from MC sources to forwarding native traffic. When the RP identifies that a new MC source started to send packets, it initiates an establishment of a native forwarding path from the DR of that source to itself. For this purpose it starts to send Join messages towards MC source to nearest neighbor to that source according the MRIB. This is a source specific Join and is noted as (S,G). When data path is established up to the DR, the DR switches from tunneling MC packets to their native forwarding, so the RP does not need to decapsulate MC packets anymore, but still continue to distribute the packets along with shared tree.

On the third phase multicast receivers will try to switch from shared tree to source specific tree by creating a direct distribution path from a multicast source. When last hop router of the multicast receiver identifies multicast traffic coming from any multicast source it will start to send Join messages towards the source with purpose to create a direct source specific path to that source. Once such path will be established and Designated router that is attached to the source L2

network will start to distribute the multicast traffic directly bypassing shared tree, the last hop router will detach its receivers from shared tree for that data and will switch to the shortest path tree distribution.

6.6.2 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.

6.6.3 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 and not VLAN interface. Connecting two routers via VLAN interface with PIM load-sharing causes loops in the network.

6.6.3.1 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.

6.6.3.2 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.

6.6.4 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.

6.6.5 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
```

6.6.5.1 Configuring IGMP

IGMP is enabled when IP multicast is enabled and static multicast or PIM is enabled on the interface.

6.6.5.2 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)#
```

6.6.5.3 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
```

6.6.6 Commands

6.6.6.1 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 Config Interface Loopback Config Interface Port Channel configured as a router port 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 Config Interface Loopback Config Interface Port Channel configured as a router port 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 Config Interface Port Channel configured as a router port
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 Config Interface Port Channel configured as a router port	
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 Config Interface Port Channel configured as a router port	
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 Config Interface Port Channel configured as a router port	
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 Config Interface Port Channel configured as a router port
History	3.3.5006
Role	admin
Example	<pre>switch (config interface vlan 10) # ip pim border</pre>
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 Config Interface Port Channel configured as a router port
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 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 switch (config) #</pre>
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: No prunes included switch (config) #</pre>
Related Commands	
Note	Should contain the following information: neighbor address, interface address, group range, Joins, Prunes.

6.6.6.2 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 Config Interface Port Channel configured as a router port	
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	
Role	admin	
Example	<pre>arc-switch14 [standalone: master] (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 Timers: Uptime/Expires Interface state: Interface, State/Mode (*, 225.0.0.0/24), 00D 00:14:49, RP 18.18.18.14, flags: BR Bidir-Upstream: Lo7 Outgoing interface list:</pre>	
Related Commands	N/A	
Note		

6.6.6.3 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 Config Interface Port Channel configured as a router port
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 Config Interface Port Channel configured as a router port	
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 Config Interface Port Channel configured as a router port	
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 Config Interface Port Channel configured as a router port	
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 Config Interface Port Channel configured as a router port	
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 Config Interface Port Channel configured as a router port	
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 Config Interface Port Channel configured as a router port	
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	<pre>switch (config)# clear ip igmp groups all switch (config)#</pre>	
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

6.7 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.

6.7.1 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 31: Common VRRP Configuration with Load Balancing



6.7.2 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.

6.7.3 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)#
```

6.7.4 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	<pre>switch (config interface vlan 10)# switch (config interface vlan 10 vrrp 10)#</pre>	
Related Commands		
Note	Maximum of 10 VRRP instances are supported.	

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	<pre>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</pre>	
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		

6.8 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.

6.8.1 MAGP Configuration

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 verify the MAGP configuration, run:**

```
switch (config)# show magp 1  
MAGP 1  
  Interface vlan:1  
  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.

6.8.2 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	<pre>switch (config)# protocol magp switch (config)#</pre>
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	switch (config interface vlan 10)# magp 1 switch (config interface vlan 10 magp 1)#
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	<pre>switch (config interface vlan 10 magp 1)# protocol magp switch (config interface vlan 10 magp 1)#</pre>
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	

6.9 DHCP Relay



DHCP Relay 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 relay agents receive messages from DHCP clients and forward them to DHCP servers thus extending the reach of the DHCP beyond the local network.

6.9.1 Commands

ip dhcp relay address

ip dhcp relay address <ip-address>
no ip dhcp relay address <ip-address>

Configures IP address of the DHCP server to forward DHCP requests.
 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	
History	3.3.4150	
Role	admin	
Example	<pre>switch (config)# ip dhcp relay address 10.10.10.10 switch (config)#</pre>	
Related Commands	N/A	
Note	- Up to 16 IP addresses may be configured - 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”	

ip dhcp relay information option

ip dhcp relay information option
no ip dhcp relay information option

Enables the DHCP relay agent to insert option 82 info on the packets.
The no form of the command removes option 82 from the packets.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config
History	3.3.4150
Role	admin
Example	<pre>switch (config)# ip dhcp relay information option switch (config)#</pre>
Related Commands	N/A
Note	

ip dhcp relay always-on

ip dhcp relay always-on
no ip dhcp relay always-on

Broadcasts DHCP requests to all interfaces with the DHCP relay agent.
 The no form of the command disables the “always-on” mode.

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config
History	3.3.4150
Role	admin
Example	<pre>switch (config)# ip dhcp relay always-on switch (config)#</pre>
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

clear ip dhcp relay counters

clear ip dhcp relay counters

Clears all DHCP relay counters (all interfaces).

Syntax Description	N/A
Default	Disabled
Configuration Mode	Config
History	3.3.4150
Role	admin
Example	<pre>switch (config)# clear ip dhcp relay counters switch (config)#</pre>
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

6.9.1.1 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 tags that the switch attaches to DHCP requests before they are forwarded.

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 configured as a router port Config Interface Port Channel configured as a router port	
History	3.3.4150	
Role	admin	
Example	<pre>switch (config interface vlan 10)# ip dhcp relay information options circuit-id my-label switch (config interface vlan 10)#</pre>	
Related Commands	N/A	
Note		

clear ip dhcp relay counters

ip dhcp relay counters
no ip dhcp relay counters

Clears all DHCP relay counters on the interface.

Syntax Description	N/A
Default	N/A
Configuration Mode	Config Interface VLAN Config Interface Ethernet configured as a router port Config Interface Port Channel configured as a router port
History	3.3.4150
Role	admin
Example	switch (config interface vlan 10)# clear ip dhcp relay counters switch (config interface vlan 10)#
Related Commands	N/A
Note	

6.9.1.2 Show

show ip dhcp relay

show ip dhcp relay

Displays DHCP relay configuration and status.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.4150
Role	admin
Example	<pre>switch (config)# show ip dhcp relay DHCP servers: 172.22.22.11, 172.33.33.33, ... (or N/A) DHCP clients requests are processed on all interfaces DHCP server responses are processed on all interfaces DHCP relay agent information option is {enabled, disabled} DHCP relay agent always-on is {enabled, disabled} Interface Label ----- Vlan10 my-label switch (config)#</pre>
Related Commands	N/A
Note	

show ip dhcp relay counters

show ip dhcp relay counters

Displays the DHCP relay counters.

Syntax Description	N/A
Default	N/A
Configuration Mode	Any Command Mode
History	3.3.4150
Role	admin
Example	<pre>switch (config)#show ip dhcp relay counters</pre> <pre> Interface Received Forwarded Dropped ----- All Req 376 376 0 All Resp 277 277 0 Interface Received Forwarded Dropped Last cleared ----- vlan1000 1000 1000 0 <Date> vlan1020 2000 2000 0 <Date> // <Date> == <DD-MM-YY, Hour-Minutes-Seconds> - something like: "20-07-13, 22:34:36" switch (config)# </pre>
Related Commands	N/A
Note	

Appendix A: Enhancing System Security According to NIST SP 800-131A

A.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.

A.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”.

A.3 Code Signing

Code signing is used to verify that the data in the image is not modified by any third-party. MLNX-OS supports signing the image files with SHA256, RSA2048 using GnuPG.



Strict mode is operational by default.

A.4 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.

A.5 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 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.

A.6 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
```

A.7 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.

➤ **To enable all encryption methods (enabled by default), run:**

```
switch (config) # ldap ssl ciphers TLS1.2
```

➤ **To verify which encryption methods are used, run:**

```
switch (config)# show ldap
User base DN : ou=People,dc=test,dc=com
User search scope : subtree
Login attribute : uid
Bind DN : cn=manager,dc=test,dc=com
Bind password : *****
Group base DN :
Group attribute : member
LDAP version : 3
Referrals : yes
Server port : 389 (not active)
Search Timeout : 5
```

```

Bind Timeout : 5
SSL mode : ssl
Server SSL port : 636
SSL ciphers : TLS1.2
SSL cert verify : yes
SSL ca-list : default-ca-list

LDAP servers:
1: 10.134.47.5
switch (config)#

```



Please make sure that “(not active)” does not appear adjacent to the line “SSL ciphers”.

A.8 Password Hashing

To comply with NIST 800-131a, Mellanox switches support password encryption with SHA512 algorithm.

➤ *To see the password encryption used, run:*

```

switch (config)# show usernames

```

USERNAME	FULL NAME	CAPABILITY	ACCOUNT STATUS
admin	System Administrator	admin	No password required for login
monitor	System Monitor	monitor	Password set (SHA512)
xmladmin	XML Admin User	admin	No password required for login
xmluser	XML Monitor User	monitor	No password required for login



Using default usernames and passwords or using usernames without passwords is highly not recommended.

When moving to strict mode, the password of each user must be reconfigured to a non-default value using the CLI command `username`.

For example, if you have a user ID “myuser” whose password is hashed with MD5, this user must be recreated manually using the command “`username myuser password mypassword`”. The password then is automatically hashed using SHA512.

The following output demonstrates the example above:

```

switch (config)# show usernames

```

USERNAME	FULL NAME	CAPABILITY	ACCOUNT STATUS
admin	System Administrator	admin	No password required for login
myuser	System Monitor	monitor	Password set (MD5)

```

switch (config)# username myuser password mypassword
switch (config)# show usernames

```

USERNAME	FULL NAME	CAPABILITY	ACCOUNT STATUS
admin	System Administrator	admin	No password required for login
myuser	System Monitor	monitor	Password set (SHA512)

Appendix B: Security Vulnerabilities and Exposures

Table 26 presents the status of common vulnerabilities and security exposures that may affect MLNX-OS.

Table 26 - Common Vulnerabilities and Exposures

CVE	Vulnerability ¹	Description	Fixed in Version
CVE-1999-0517	Not vulnerable: Requires strict mode	An SNMP community name is the default (e.g. public), null, or missing.	3.4.0000
CVE-2006-0175	Not vulnerable	Directory traversal vulnerability in scp for OpenSSH before 3.4p1 allows remote malicious servers to overwrite arbitrary files. NOTE: this may be a rediscovery of CVE-2000-0992.	3.3.3500
CVE-2006-1653	N/A	The default configuration for OpenSSH enables AllowTcpForwarding, which could allow remote authenticated users to perform a port bounce, when configured with an anonymous access program such as AnonCVS.	N/A
CVE-2006-2760	N/A	sshd in OpenSSH 3.5p1, when PermitRootLogin is disabled, immediately closes the TCP connection after a root login attempt with the correct password, but leaves the connection open after an attempt with an incorrect password, which makes it easier for remote attackers to guess the password by observing the connection state, a different vulnerability than CVE-2003-0190. NOTE: it could be argued that in most environments, this does not cross privilege boundaries without requiring leverage of a separate vulnerability.	N/A
CVE-2005-2797	N/A	OpenSSH 4.0, and other versions before 4.2, does not properly handle dynamic port forwarding ("-D" option) when a listen address is not provided, which may cause OpenSSH to enable the GatewayPorts functionality.	N/A
CVE-2005-2798	Not vulnerable	sshd in OpenSSH before 4.2, when GSSAPIDelegateCredentials is enabled, allows GSSAPI credentials to be delegated to clients who log in using non-GSSAPI methods, which could cause those credentials to be exposed to untrusted users or hosts.	3.3.3500
CVE-2006-0225	N/A	scp in OpenSSH 4.2p1 allows attackers to execute arbitrary commands via filenames that contain shell metacharacters or spaces, which are expanded twice.	N/A
CVE-2006-4924	Not vulnerable	sshd in OpenSSH before 4.4, when using the version 1 SSH protocol, allows remote attackers to cause a denial of service (CPU consumption) via an SSH packet that contains duplicate blocks, which is not properly handled by the CRC compensation attack detector.	3.4.0000

Table 26 - Common Vulnerabilities and Exposures

CVE	Vulnerability ¹	Description	Fixed in Version
CVE-2006-4925	N/A	packet.c in ssh in OpenSSH allows remote attackers to cause a denial of service (crash) by sending an invalid protocol sequence with USERAUTH_SUCCESS before NEWKEYS, which causes newkeys[mode] to be NULL.	N/A
CVE-2006-5051	N/A	Signal handler race condition in OpenSSH before 4.4 allows remote attackers to cause a denial of service (crash), and possibly execute arbitrary code if GSSAPI authentication is enabled, via unspecified vectors that lead to a double-free.	N/A
CVE-2006-5052	N/A	Unspecified vulnerability in portable OpenSSH before 4.4, when running on some platforms, allows remote attackers to determine the validity of usernames via unknown vectors involving a GSSAPI "authentication abort."	N/A
CVE-2006-5229	N/A	OpenSSH portable 4.1 on SUSE Linux, and possibly other platforms and versions, and possibly under limited configurations, allows remote attackers to determine valid usernames via timing discrepancies in which responses take longer for valid usernames than invalid ones, as demonstrated by sstime. NOTE: as of 20061014, it appears that this issue is dependent on the use of manually-set passwords that causes delays when processing /etc/shadow due to an increased number of rounds.	N/A
CVE-2006-5794	Not vulnerable	Unspecified vulnerability in the sshd Privilege Separation Monitor in OpenSSH before 4.5 causes weaker verification that authentication has been successful, which might allow attackers to bypass authentication. NOTE: as of 20061108, it is believed that this issue is only exploitable by leveraging vulnerabilities in the unprivileged process, which are not known to exist.	3.4.0000
CVE-2007-0726	N/A	The SSH key generation process in OpenSSH in Apple Mac OS X 10.3.9 and 10.4 through 10.4.8 allows remote attackers to cause a denial of service by connecting to the server before SSH has finished creating keys, which causes the keys to be regenerated and can break trust relationships that were based on the original keys.	N/A
CVE-2007-2243	N/A	OpenSSH 4.6 and earlier, when ChallengeResponseAuthentication is enabled, allows remote attackers to determine the existence of user accounts by attempting to authenticate via S/KEY, which displays a different response if the user account exists, a similar issue to CVE-2001-1483.	N/A

Table 26 - Common Vulnerabilities and Exposures

CVE	Vulnerability ¹	Description	Fixed in Version
CVE-2007-2768	N/A	OpenSSH, when using OPIE (One-Time Passwords in Everything) for PAM, allows remote attackers to determine the existence of certain user accounts, which displays a different response if the user account exists and is configured to use one-time passwords (OTP), a similar issue to CVE-2007-2243.	N/A
CVE-2007-3102	N/A	Unspecified vulnerability in the <code>linux_audit_record_event</code> function in OpenSSH 4.3p2, as used on Fedora Core 6 and possibly other systems, allows remote attackers to write arbitrary characters to an audit log via a crafted username. NOTE: some of these details are obtained from third party information.	N/A
CVE-2007-4654	N/A	Unspecified vulnerability in SSHield 1.6.1 with OpenSSH 3.0.2p1 on Cisco WebNS 8.20.0.1 on Cisco Content Services Switch (CSS) series 11000 devices allows remote attackers to cause a denial of service (connection slot exhaustion and device crash) via a series of large packets designed to exploit the SSH CRC32 attack detection overflow (CVE-2001-0144), possibly a related issue to CVE-2002-1024.	N/A
CVE-2007-4752	Not vulnerable	ssh in OpenSSH before 4.7 does not properly handle when an untrusted cookie cannot be created and uses a trusted X11 cookie instead, which allows attackers to violate intended policy and gain privileges by causing an X client to be treated as trusted.	3.4.0000
CVE-2007-5715	N/A	DenyHosts 2.6 processes OpenSSH sshd "not listed in AllowUsers" log messages with an incorrect regular expression that does not match an IP address, which might allow remote attackers to avoid detection and blocking when making invalid login attempts with a username not present in AllowUsers, as demonstrated by the root username, a different vulnerability than CVE-2007-4323.	N/A
CVE-2007-6415	N/A	scponly 4.6 and earlier allows remote authenticated users to bypass intended restrictions and execute arbitrary code by invoking scp, as implemented by OpenSSH, with the -F and -o options.	N/A
CVE-2008-1483	Not vulnerable	OpenSSH 4.3p2, and probably other versions, allows local users to hijack forwarded X connections by causing ssh to set DISPLAY to :10, even when another process is listening on the associated port, as demonstrated by opening TCP port 6010 (IPv4) and sniffing a cookie sent by Emacs.	3.4.0000
CVE-2008-1657	N/A	OpenSSH 4.4 up to versions before 4.9 allows remote authenticated users to bypass the sshd_config ForceCommand directive by modifying the .ssh/rc session file.	N/A

Table 26 - Common Vulnerabilities and Exposures

CVE	Vulnerability ¹	Description	Fixed in Version
CVE-2008-3234	Only affects OpenSSH version 4.x. MLNX-OS uses 3.8.1p1.	sshd in OpenSSH 4 on Debian GNU/Linux, and the 20070303 OpenSSH snapshot, allows remote authenticated users to obtain access to arbitrary SELinux roles by appending a :/ (colon slash) sequence, followed by the role name, to the username.	N/A
CVE-2008-3259	N/A	OpenSSH before 5.1 sets the SO_REUSEADDR socket option when the X11UseLocalhost configuration setting is disabled, which allows local users on some platforms to hijack the X11 forwarding port via a bind to a single IP address, as demonstrated on the HP-UX platform.	N/A
CVE-2008-3844	N/A	Certain Red Hat Enterprise Linux (RHEL) 4 and 5 packages for OpenSSH, as signed in August 2008 using a legitimate Red Hat GPG key, contain an externally introduced modification (Trojan Horse) that allows the package authors to have an unknown impact. NOTE: since the malicious packages were not distributed from any official Red Hat sources, the scope of this issue is restricted to users who may have obtained these packages through unofficial distribution points. As of 20080827, no unofficial distributions of this software are known.	N/A
CVE-2008-4109	Not vulnerable	A certain Debian patch for OpenSSH before 4.3p2-9etch3 on etch; before 4.6p1-1 on sid and lenny; and on other distributions such as SUSE uses functions that are not async-signal-safe in the signal handler for login timeouts, which allows remote attackers to cause a denial of service (connection slot exhaustion) via multiple login attempts. NOTE: this issue exists because of an incorrect fix for CVE-2006-5051.	3.4.0000
CVE-2008-5161	Not vulnerable	Error handling in the SSH protocol in (1) SSH Tectia Client and Server and Connector 4.0 through 4.4.11, 5.0 through 5.2.4, and 5.3 through 5.3.8; Client and Server and ConnectSecure 6.0 through 6.0.4; Server for Linux on IBM System z 6.0.4; Server for IBM z/OS 5.5.1 and earlier, 6.0.0, and 6.0.1; and Client 4.0-J through 4.3.3-J and 4.0-K through 4.3.10-K; and (2) OpenSSH 4.7p1 and possibly other versions, when using a block cipher algorithm in Cipher Block Chaining (CBC) mode, makes it easier for remote attackers to recover certain plaintext data from an arbitrary block of ciphertext in an SSH session via unknown vectors.	3.4.0000

Table 26 - Common Vulnerabilities and Exposures

CVE	Vulnerability ¹	Description	Fixed in Version
CVE-2009-2904	N/A	A certain Red Hat modification to the ChrootDirectory feature in OpenSSH 4.8, as used in sshd in OpenSSH 4.3 in Red Hat Enterprise Linux (RHEL) 5.4 and Fedora 11, allows local users to gain privileges via hard links to setuid programs that use configuration files within the chroot directory, related to requirements for directory ownership.	N/A
CVE-2010-4478	N/A	OpenSSH 5.6 and earlier, when J-PAKE is enabled, does not properly validate the public parameters in the J-PAKE protocol, which allows remote attackers to bypass the need for knowledge of the shared secret, and successfully authenticate, by sending crafted values in each round of the protocol, a related issue to CVE-2010-4252.	N/A
CVE-2010-4755	N/A	The (1) remote_glob function in sftp-glob.c and the (2) process_put function in sftp.c in OpenSSH 5.8 and earlier, as used in FreeBSD 7.3 and 8.1, NetBSD 5.0.2, OpenBSD 4.7, and other products, allow remote authenticated users to cause a denial of service (CPU and memory consumption) via crafted glob expressions that do not match any pathnames, as demonstrated by glob expressions in SSH_FXP_STAT requests to an sftp daemon, a different vulnerability than CVE-2010-2632.	N/A
CVE-2010-5298	Not vulnerable	Race condition in the ssl3_read_bytes function in s3_pkt.c in OpenSSL through 1.0.1g, when SSL_MODE_RELEASE_BUFFERS is enabled, allows remote attackers to inject data across sessions or cause a denial of service (use-after-free and parsing error) via an SSL connection in a multithreaded environment.	3.4.0008
CVE-2011-0539	N/A	The key_certify function in usr.bin/ssh/key.c in OpenSSH 5.6 and 5.7, when generating legacy certificates using the -t command-line option in ssh-keygen, does not initialize the nonce field, which might allow remote attackers to obtain sensitive stack memory contents or make it easier to conduct hash collision attacks.	N/A
CVE-2011-3389	Not vulnerable	The SSL protocol, as used in certain configurations in Microsoft Windows and Microsoft Internet Explorer, Mozilla Firefox, Google Chrome, Opera, and other products, encrypts data by using CBC mode with chained initialization vectors, which allows man-in-the-middle attackers to obtain plaintext HTTP headers via a blockwise chosen-boundary attack (BCBA) on an HTTPS session, in conjunction with JavaScript code that uses (1) the HTML5 WebSocket API, (2) the Java URLConnection API, or (3) the Silverlight WebClient API, aka a "BEAST" attack.	3.3.3500

Table 26 - Common Vulnerabilities and Exposures

CVE	Vulnerability ¹	Description	Fixed in Version
CVE-2011-3607	Not vulnerable	Integer overflow in the ap_pregsub function in server/util.c in the Apache HTTP Server 2.0.x through 2.0.64 and 2.2.x through 2.2.21, when the mod_setenvif module is enabled, allows local users to gain privileges via a .htaccess file with a crafted SetEnvIf directive, in conjunction with a crafted HTTP request header, leading to a heap-based buffer overflow.	3.3.3500
CVE-2011-4317	Not vulnerable	The mod_proxy module in the Apache HTTP Server 1.3.x through 1.3.42, 2.0.x through 2.0.64, and 2.2.x through 2.2.21, when the Revision 1179239 patch is in place, does not properly interact with use of (1) RewriteRule and (2) ProxyPassMatch pattern matches for configuration of a reverse proxy, which allows remote attackers to send requests to intranet servers via a malformed URI containing an @ (at sign) character and a : (colon) character in invalid positions. NOTE: this vulnerability exists because of an incomplete fix for CVE-2011-3368.	3.3.3500
CVE-2011-4327	N/A	ssh-keysign.c in ssh-keysign in OpenSSH before 5.8p2 on certain platforms executes ssh-rand-helper with unintended open file descriptors, which allows local users to obtain sensitive key information via the ptrace system call.	N/A
CVE-2011-5000	N/A	The ssh_gssapi_parse_ename function in gss-serv.c in OpenSSH 5.8 and earlier, when gssapi-with-mic authentication is enabled, allows remote authenticated users to cause a denial of service (memory consumption) via a large value in a certain length field. NOTE: there may be limited scenarios in which this issue is relevant.	N/A
CVE-2012-0031	Not vulnerable	scoreboard.c in the Apache HTTP Server 2.2.21 and earlier might allow local users to cause a denial of service (daemon crash during shutdown) or possibly have unspecified other impact by modifying a certain type field within a scoreboard shared memory segment, leading to an invalid call to the free function.	3.3.3500
CVE-2012-0053	Not vulnerable	protocol.c in the Apache HTTP Server 2.2.x through 2.2.21 does not properly restrict header information during construction of Bad Request (aka 400) error documents, which allows remote attackers to obtain the values of HTTPOnly cookies via vectors involving a (1) long or (2) malformed header in conjunction with crafted web script.	3.3.3500

Table 26 - Common Vulnerabilities and Exposures

CVE	Vulnerability ¹	Description	Fixed in Version
CVE-2012-0814	Not vulnerable	The <code>auth_parse_options</code> function in <code>auth-options.c</code> in <code>sshd</code> in OpenSSH before 5.7 provides debug messages containing <code>authorized_keys</code> command options, which allows remote authenticated users to obtain potentially sensitive information by reading these messages, as demonstrated by the shared user account required by Gitolite. NOTE: this can cross privilege boundaries because a user account may intentionally have no shell or filesystem access, and therefore may have no supported way to read an <code>authorized_keys</code> file in its own home directory.	3.3.3500
CVE-2012-2687	N/A	Multiple cross-site scripting (XSS) vulnerabilities in the <code>make_variant_list</code> function in <code>mod_negotiation.c</code> in the <code>mod_negotiation</code> module in the Apache HTTP Server 2.4.x before 2.4.3, when the MultiViews option is enabled, allow remote attackers to inject arbitrary web script or HTML via a crafted filename that is not properly handled during construction of a variant list.	N/A
CVE-2012-4929	Not vulnerable	The TLS protocol 1.2 and earlier, as used in Mozilla Firefox, Google Chrome, Qt, and other products, can encrypt compressed data without properly obfuscating the length of the unencrypted data, which allows man-in-the-middle attackers to obtain plaintext HTTP headers by observing length differences during a series of guesses in which a string in an HTTP request potentially matches an unknown string in an HTTP header, aka a "CRIME" attack.	3.3.3500
CVE-2013-7423	Not vulnerable	The <code>send_dg</code> function in <code>resolv/res_send.c</code> in GNU C Library (aka glibc or libc6) before 2.20 does not properly reuse file descriptors, which allows remote attackers to send DNS queries to unintended locations via a large number of request that trigger a call to the <code>getaddrinfo</code> function.	3.4.3000
CVE-2014-0195	Not vulnerable	The <code>dtls1_reassemble_fragment</code> function in <code>d1_both.c</code> in OpenSSL before 0.9.8za, 1.0.0 before 1.0.0m, and 1.0.1 before 1.0.1h does not properly validate fragment lengths in DTLS ClientHello messages, which allows remote attackers to execute arbitrary code or cause a denial of service (buffer overflow and application crash) via a long non-initial fragment.	3.4.0008
CVE-2014-0198	Not vulnerable	The <code>do_ssl3_write</code> function in <code>s3_pkt.c</code> in OpenSSL 1.x through 1.0.1g, when <code>SSL_MODE_RELEASE_BUFFERS</code> is enabled, does not properly manage a buffer pointer during certain recursive calls, which allows remote attackers to cause a denial of service (NULL pointer dereference and application crash) via vectors that trigger an alert condition.	3.4.0008

Table 26 - Common Vulnerabilities and Exposures

CVE	Vulnerability ¹	Description	Fixed in Version
CVE-2014-0221	Not vulnerable	The <code>dtls1_get_message_fragment</code> function in <code>d1_both.c</code> in OpenSSL before 0.9.8za, 1.0.0 before 1.0.0m, and 1.0.1 before 1.0.1h allows remote attackers to cause a denial of service (recursion and client crash) via a DTLS hello message in an invalid DTLS handshake.	3.4.0008
CVE-2014-0224	Not vulnerable	OpenSSL before 0.9.8za, 1.0.0 before 1.0.0m, and 1.0.1 before 1.0.1h does not properly restrict processing of <code>ChangeCipherSpec</code> messages, which allows man-in-the-middle attackers to trigger use of a zero-length master key in certain OpenSSL-to-OpenSSL communications, and consequently hijack sessions or obtain sensitive information, via a crafted TLS handshake, aka the "CCS Injection" vulnerability.	3.4.1000
CVE-2014-0475	Not vulnerable	Multiple directory traversal vulnerabilities in GNU C Library (aka glibc or libc6) before 2.20 allow context-dependent attackers to bypass <code>ForceCommand</code> restrictions and possibly have other unspecified impact via a <code>..</code> (dot dot) in a (1) <code>LC_*</code> , (2) <code>LANG</code> , or other locale environment variable.	3.4.3000
CVE-2014-1692	N/A	The <code>hash_buffer</code> function in <code>schorr.c</code> in OpenSSH through 6.4, when <code>Makefile.inc</code> is modified to enable the J-PAKE protocol, does not initialize certain data structures, which might allow remote attackers to cause a denial of service (memory corruption) or have unspecified other impact via vectors that trigger an error condition.	N/A
CVE-2014-2532	N/A	<code>sshd</code> in OpenSSH before 6.6 does not properly support wildcards on <code>AcceptEnv</code> lines in <code>sshd_config</code> , which allows remote attackers to bypass intended environment restrictions by using a substring located before a wildcard character.	N/A
CVE-2014-3470	Not vulnerable	The <code>ssl3_send_client_key_exchange</code> function in <code>s3_clnt.c</code> in OpenSSL before 0.9.8za, 1.0.0 before 1.0.0m, and 1.0.1 before 1.0.1h, when an anonymous ECDH cipher suite is used, allows remote attackers to cause a denial of service (NULL pointer dereference and client crash) by triggering a NULL certificate value.	3.4.0008
CVE-2014-3505	Not vulnerable	Double free vulnerability in <code>d1_both.c</code> in the DTLS implementation in OpenSSL 0.9.8 before 0.9.8zb, 1.0.0 before 1.0.0n, and 1.0.1 before 1.0.1i allows remote attackers to cause a denial of service (application crash) via crafted DTLS packets that trigger an error condition.	3.4.1000

Table 26 - Common Vulnerabilities and Exposures

CVE	Vulnerability ¹	Description	Fixed in Version
CVE-2014-3506	Not vulnerable	d1_both.c in the DTLS implementation in OpenSSL 0.9.8 before 0.9.8zb, 1.0.0 before 1.0.0n, and 1.0.1 before 1.0.1i allows remote attackers to cause a denial of service (memory consumption) via crafted DTLS handshake messages that trigger memory allocations corresponding to large length values.	3.4.1000
CVE-2014-3507	Not vulnerable	Memory leak in d1_both.c in the DTLS implementation in OpenSSL 0.9.8 before 0.9.8zb, 1.0.0 before 1.0.0n, and 1.0.1 before 1.0.1i allows remote attackers to cause a denial of service (memory consumption) via zero-length DTLS fragments that trigger improper handling of the return value of a certain insert function.	3.4.1000
CVE-2014-3508	Not vulnerable	The OBJ_obj2txt function in crypto/objects/obj_dat.c in OpenSSL 0.9.8 before 0.9.8zb, 1.0.0 before 1.0.0n, and 1.0.1 before 1.0.1i, when pretty printing is used, does not ensure the presence of '\0' characters, which allows context-dependent attackers to obtain sensitive information from process stack memory by reading output from X509_name_online, X509_name_print_ex, and unspecified other functions.	3.4.1000
CVE-2014-3509	Not vulnerable	Race condition in the ssl_parse_serverhello_tlsext function in tl_lib.c in OpenSSL 1.0.0 before 1.0.0n and 1.0.1 before 1.0.1i, when multithreading and session resumption are used, allows remote SSL servers to cause a denial of service (memory overwrite and client application crash) or possibly have unspecified other impact by sending Elliptic Curve (EC) Supported Point Formats Extension data.	3.4.1000
CVE-2014-3510	Not vulnerable	The ssl3_send_client_key_exchange function in s3_clnt.c in OpenSSL 0.9.8 before 0.9.8zb, 1.0.0 before 1.0.0n, and 1.0.1 before 1.0.1i allows remote DTLS servers to cause a denial of service (NULL pointer dereference and client application crash) via a crafted handshake message in conjunction with a (1) anonymous DH or (2) anonymous ECDH ciphersuite.	3.4.1000
CVE-2014-3511	Not vulnerable	The ssl23_get_client_hello function in s23_srvr.c in OpenSSL 1.0.1 before 1.0.1i allows man-in-the-middle attackers to force the use of TLS 1.0 by triggering ClientHello message fragmentation in communication between a client and server that both support later TLS versions, related to a "protocol downgrade" issue.	3.4.1000
CVE-2014-3513	Not vulnerable	Memory leak in d1_srtp.c in the DTLS SRTP extension in OpenSSL 1.0.1 before 1.0.1j allows remote attackers to cause a denial of service (memory consumption) via a crafted handshake message.	3.4.1000

Table 26 - Common Vulnerabilities and Exposures

CVE	Vulnerability ¹	Description	Fixed in Version
CVE-2014-3566	Not vulnerable	The SSL protocol 3.0, as used in OpenSSL through 1.0.1i and other products, uses nondeterministic CBC padding, which makes it easier for man-in-the-middle attackers to obtain cleartext data via a padding-oracle attack, aka the "POODLE" issue.	3.4.1000
CVE-2014-3567	Not vulnerable	Memory leak in the <code>tls_decrypt_ticket</code> function in <code>t1_lib.c</code> in OpenSSL before 0.9.8zc, 1.0.0 before 1.0.0o, and 1.0.1 before 1.0.1j allows remote attackers to cause a denial of service (memory consumption) via a crafted session ticket that triggers an integrity-check failure.	3.4.1000
CVE-2014-3569	N/A	The <code>ssl23_get_client_hello</code> function in <code>s23_srvr.c</code> in OpenSSL 0.9.8zc, 1.0.0o, and 1.0.1j does not properly handle attempts to use unsupported protocols, which allows remote attackers to cause a denial of service (NULL pointer dereference and daemon crash) via an unexpected handshake, as demonstrated by an SSLv3 handshake to a no-ssl3 application with certain error handling. NOTE: this issue became relevant after the CVE-2014-3568 fix.	N/A
CVE-2014-3570	Not vulnerable	The <code>BN_sqr</code> implementation in OpenSSL before 0.9.8zd, 1.0.0 before 1.0.0p, and 1.0.1 before 1.0.1k does not properly calculate the square of a BIGNUM value, which might make it easier for remote attackers to defeat cryptographic protection mechanisms via unspecified vectors, related to <code>crypto/bn/asm/mips.pl</code> , <code>crypto/bn/asm/x86_64-gcc.c</code> , and <code>crypto/bn/bn_asm.c</code> .	3.4.3000
CVE-2014-3571	Not vulnerable	OpenSSL before 0.9.8zd, 1.0.0 before 1.0.0p, and 1.0.1 before 1.0.1k allows remote attackers to cause a denial of service (NULL pointer dereference and application crash) via a crafted DTLS message that is processed with a different read operation for the handshake header than for the handshake body, related to the <code>dtls1_get_record</code> function in <code>d1_pkt.c</code> and the <code>ssl3_read_n</code> function in <code>s3_pkt.c</code> .	3.4.3000
CVE-2014-3572	Not vulnerable	The <code>ssl3_get_key_exchange</code> function in <code>s3_clnt.c</code> in OpenSSL before 0.9.8zd, 1.0.0 before 1.0.0p, and 1.0.1 before 1.0.1k allows remote SSL servers to conduct ECDHE-to-ECDH downgrade attacks and trigger a loss of forward secrecy by omitting the <code>ServerKeyExchange</code> message.	3.4.3000
CVE-2014-6040	Not vulnerable	GNU C Library (aka glibc) before 2.20 allows context-dependent attackers to cause a denial of service (out-of-bounds read and crash) via a multibyte character value of "0xffff" to the <code>iconv</code> function when converting (1) IBM933, (2) IBM935, (3) IBM937, (4) IBM939, or (5) IBM1364 encoded data to UTF-8.	3.4.2300

Table 26 - Common Vulnerabilities and Exposures

CVE	Vulnerability ¹	Description	Fixed in Version
CVE-2014-6271	Not vulnerable	GNU Bash through 4.3 processes trailing strings after function definitions in the values of environment variables, which allows remote attackers to execute arbitrary code via a crafted environment, as demonstrated by vectors involving the ForceCommand feature in OpenSSH sshd, the mod_cgi and mod_cgid modules in the Apache HTTP Server, scripts executed by unspecified DHCP clients, and other situations in which setting the environment occurs across a privilege boundary from Bash execution, aka "ShellShock." NOTE: the original fix for this issue was incorrect; CVE-2014-7169 has been assigned to cover the vulnerability that is still present after the incorrect fix.	3.4.0008
CVE-2014-7817	Not vulnerable	The wordexp function in GNU C Library (aka glibc) 2.21 does not enforce the WRDE_NOCMD flag, which allows context-dependent attackers to execute arbitrary commands, as demonstrated by input containing "\$((`...`))".	3.4.3000
CVE-2014-8176	Not vulnerable	The dtls1_clear_queues function in ssl/d1_lib.c in OpenSSL before 0.9.8za, 1.0.0 before 1.0.0m, and 1.0.1 before 1.0.1h frees data structures without considering that application data can arrive between a ChangeCipherSpec message and a Finished message, which allows remote DTLS peers to cause a denial of service (memory corruption and application crash) or possibly have unspecified other impact via unexpected application data.	3.4.3000
CVE-2014-8275	Not vulnerable	OpenSSL before 0.9.8zd, 1.0.0 before 1.0.0p, and 1.0.1 before 1.0.1k does not enforce certain constraints on certificate data, which allows remote attackers to defeat a fingerprint-based certificate-blacklist protection mechanism by including crafted data within a certificate's unsigned portion, related to crypto/asn1/a_verify.c, crypto/dsa/dsa_asn1.c, crypto/ecdsa/ecdsa_vrf.c, and crypto/x509/x_all.c.	3.4.3000
CVE-2014-9293	Not vulnerable	The config_auth function in ntpd in NTP before 4.2.7p11, when an auth key is not configured, improperly generates a key, which makes it easier for remote attackers to defeat cryptographic protection mechanisms via a brute-force attack.	3.4.1000
CVE-2014-9294	Not vulnerable	util/ntp-keygen.c in ntp-keygen in NTP before 4.2.7p230 uses a weak RNG seed, which makes it easier for remote attackers to defeat cryptographic protection mechanisms via a brute-force attack.	3.4.1000

Table 26 - Common Vulnerabilities and Exposures

CVE	Vulnerability ¹	Description	Fixed in Version
CVE-2014-9295	Not vulnerable	Multiple stack-based buffer overflows in ntpd in NTP before 4.2.8 allow remote attackers to execute arbitrary code via a crafted packet, related to (1) the crypto_rcv function when the Autokey Authentication feature is used, (2) the ctl_putdata function, and (3) the configure function.	3.4.1000
CVE-2014-9296	Not vulnerable	The receive function in ntp_proto.c in ntpd in NTP before 4.2.8 continues to execute after detecting a certain authentication error, which might allow remote attackers to trigger an unintended association change via crafted packets.	3.4.1000
CVE-2014-9297	N/A	This candidate has been reserved by an organization or individual that will use it when announcing a new security problem. When the candidate has been publicized, the details for this candidate will be provided	3.4.3000
CVE-2015-0204	Not vulnerable	The ssl3_get_key_exchange function in s3_clnt.c in OpenSSL before 0.9.8zd, 1.0.0 before 1.0.0p, and 1.0.1 before 1.0.1k allows remote SSL servers to conduct RSA-to-EXPORT_RSA downgrade attacks and facilitate brute-force decryption by offering a weak ephemeral RSA key in a noncompliant role, related to the "FREAK" issue. NOTE: the scope of this CVE is only client code based on OpenSSL, not EXPORT_RSA issues associated with servers or other TLS implementations.	3.4.3000
CVE-2015-0205	Not vulnerable	The ssl3_get_cert_verify function in s3_srvr.c in OpenSSL 1.0.0 before 1.0.0p and 1.0.1 before 1.0.1k accepts client authentication with a Diffie-Hellman (DH) certificate without requiring a CertificateVerify message, which allows remote attackers to obtain access without knowledge of a private key via crafted TLS Handshake Protocol traffic to a server that recognizes a Certification Authority with DH support.	3.4.3000
CVE-2015-0206	Not vulnerable	Memory leak in the dtls1_buffer_record function in dl_pkt.c in OpenSSL 1.0.0 before 1.0.0p and 1.0.1 before 1.0.1k allows remote attackers to cause a denial of service (memory consumption) by sending many duplicate records for the next epoch, leading to failure of replay detection.	3.4.3000
CVE-2015-0207	Not vulnerable	The dtls1_listen function in dl_lib.c in OpenSSL 1.0.2 before 1.0.2a does not properly isolate the state information of independent data streams, which allows remote attackers to cause a denial of service (application crash) via crafted DTLS traffic, as demonstrated by DTLS 1.0 traffic to a DTLS 1.2 server.	3.4.2008

Table 26 - Common Vulnerabilities and Exposures

CVE	Vulnerability ¹	Description	Fixed in Version
CVE-2015-0208	Not vulnerable	The ASN.1 signature-verification implementation in the <code>rsa_item_verify</code> function in <code>crypto/rsa/rsa_ameth.c</code> in OpenSSL 1.0.2 before 1.0.2a allows remote attackers to cause a denial of service (NULL pointer dereference and application crash) via crafted RSA PSS parameters to an endpoint that uses the certificate-verification feature.	3.4.2008
CVE-2015-0209	N/A	Use-after-free vulnerability in the <code>d2i_ECPrivateKey</code> function in <code>crypto/ec/ec_asn1.c</code> in OpenSSL before 0.9.8zf, 1.0.0 before 1.0.0r, 1.0.1 before 1.0.1m, and 1.0.2 before 1.0.2a might allow remote attackers to cause a denial of service (memory corruption and application crash) or possibly have unspecified other impact via a malformed Elliptic Curve (EC) private-key file that is improperly handled during import.	3.4.3000
CVE-2015-0235	Not vulnerable	Heap-based buffer overflow in the <code>__nss_hostname_digits_dots</code> function in <code>glibc 2.2</code> , and other 2.x versions before 2.18, allows context-dependent attackers to execute arbitrary code via vectors related to the (1) <code>gethostbyname</code> or (2) <code>gethostbyname2</code> function, aka "GHOST."	3.4.2002
CVE-2015-0285	N/A	The <code>ssl3_client_hello</code> function in <code>s3_clnt.c</code> in OpenSSL 1.0.2 before 1.0.2a does not ensure that the PRNG is seeded before proceeding with a handshake, which makes it easier for remote attackers to defeat cryptographic protection mechanisms by sniffing the network and then conducting a brute-force attack.	3.4.3000
CVE-2015-0286	Not vulnerable	The <code>ASN1_TYPE_cmp</code> function in <code>crypto/asn1/a_type.c</code> in OpenSSL before 0.9.8zf, 1.0.0 before 1.0.0r, 1.0.1 before 1.0.1m, and 1.0.2 before 1.0.2a does not properly perform boolean-type comparisons, which allows remote attackers to cause a denial of service (invalid read operation and application crash) via a crafted X.509 certificate to an endpoint that uses the certificate-verification feature.	3.4.3000
CVE-2015-0287	Not vulnerable	The <code>ASN1_item_ex_d2i</code> function in <code>crypto/asn1/tasn_dec.c</code> in OpenSSL before 0.9.8zf, 1.0.0 before 1.0.0r, 1.0.1 before 1.0.1m, and 1.0.2 before 1.0.2a does not reinitialize CHOICE and ADB data structures, which might allow attackers to cause a denial of service (invalid write operation and memory corruption) by leveraging an application that relies on ASN.1 structure reuse.	3.4.3000
CVE-2015-0288	Not vulnerable	The <code>X509_to_X509_REQ</code> function in <code>crypto/x509/x509_req.c</code> in OpenSSL before 0.9.8zf, 1.0.0 before 1.0.0r, 1.0.1 before 1.0.1m, and 1.0.2 before 1.0.2a might allow attackers to cause a denial of service (NULL pointer dereference and application crash) via an invalid certificate key.	3.4.3000

Table 26 - Common Vulnerabilities and Exposures

CVE	Vulnerability ¹	Description	Fixed in Version
CVE-2015-0289	Not vulnerable	The PKCS#7 implementation in OpenSSL before 0.9.8zf, 1.0.0 before 1.0.0r, 1.0.1 before 1.0.1m, and 1.0.2 before 1.0.2a does not properly handle a lack of outer ContentInfo, which allows attackers to cause a denial of service (NULL pointer dereference and application crash) by leveraging an application that processes arbitrary PKCS#7 data and providing malformed data with ASN.1 encoding, related to crypto/pkcs7/pk7_doit.c and crypto/pkcs7/pk7_lib.c.	3.4.3000
CVE-2015-0290	Not vulnerable	The multi-block feature in the ssl3_write_bytes function in s3_pkt.c in OpenSSL 1.0.2 before 1.0.2a on 64-bit x86 platforms with AES NI support does not properly handle certain non-blocking I/O cases, which allows remote attackers to cause a denial of service (pointer corruption and application crash) via unspecified vectors.	3.4.2008
CVE-2015-0291	Not vulnerable	The sigalgs implementation in tl1_lib.c in OpenSSL 1.0.2 before 1.0.2a allows remote attackers to cause a denial of service (NULL pointer dereference and daemon crash) by using an invalid signature_algorithms extension in the ClientHello message during a renegotiation.	3.4.2008
CVE-2015-0292	Not vulnerable	Integer underflow in the EVP_DecodeUpdate function in crypto/evp/encode.c in the base64-decoding implementation in OpenSSL before 0.9.8za, 1.0.0 before 1.0.0m, and 1.0.1 before 1.0.1h allows remote attackers to cause a denial of service (memory corruption) or possibly have unspecified other impact via crafted base64 data that triggers a buffer overflow.	3.4.3000
CVE-2015-0293	Not vulnerable	The SSLv2 implementation in OpenSSL before 0.9.8zf, 1.0.0 before 1.0.0r, 1.0.1 before 1.0.1m, and 1.0.2 before 1.0.2a allows remote attackers to cause a denial of service (s2_lib.c assertion failure and daemon exit) via a crafted CLIENT-MASTER-KEY message.	3.4.3000
CVE-2015-1787	Not vulnerable	The ssl3_get_client_key_exchange function in s3_srvr.c in OpenSSL 1.0.2 before 1.0.2a, when client authentication and an ephemeral Diffie-Hellman ciphersuite are enabled, allows remote attackers to cause a denial of service (daemon crash) via a ClientKeyExchange message with a length of zero.	3.4.2008
CVE-2015-1789	Not vulnerable	The X509_cmp_time function in crypto/x509/x509_vfy.c in OpenSSL before 0.9.8zg, 1.0.0 before 1.0.0s, 1.0.1 before 1.0.1n, and 1.0.2 before 1.0.2b allows remote attackers to cause a denial of service (out-of-bounds read and application crash) via a crafted length field in ASN1_TIME data, as demonstrated by an attack against a server that supports client authentication with a custom verification callback.	3.4.3000

Table 26 - Common Vulnerabilities and Exposures

CVE	Vulnerability ¹	Description	Fixed in Version
CVE-2015-1790	Not vulnerable	The PKCS7_dataDecodefunction in crypto/pkcs7/pk7_doit.c in OpenSSL before 0.9.8zg, 1.0.0 before 1.0.0s, 1.0.1 before 1.0.1n, and 1.0.2 before 1.0.2b allows remote attackers to cause a denial of service (NULL pointer dereference and application crash) via a PKCS#7 blob that uses ASN.1 encoding and lacks inner EncryptedContent data.	3.4.3000
CVE-2015-1791	Not vulnerable	Race condition in the ssl3_get_new_session_ticket function in ssl/s3_clnt.c in OpenSSL before 0.9.8zg, 1.0.0 before 1.0.0s, 1.0.1 before 1.0.1n, and 1.0.2 before 1.0.2b, when used for a multi-threaded client, allows remote attackers to cause a denial of service (double free and application crash) or possibly have unspecified other impact by providing a NewSessionTicket during an attempt to reuse a ticket that had been obtained earlier.	3.4.3000
CVE-2015-1792	Not vulnerable	The do_free_upto function in crypto/cms/cms_smime.c in OpenSSL before 0.9.8zg, 1.0.0 before 1.0.0s, 1.0.1 before 1.0.1n, and 1.0.2 before 1.0.2b allows remote attackers to cause a denial of service (infinite loop) via vectors that trigger a NULL value of a BIO data structure, as demonstrated by an unrecognized X.660 OID for a hash function.	3.4.3000
CVE-2015-1798	N/A	The symmetric-key feature in the receive function in ntp_proto.c in ntpd in NTP 4.x before 4.2.8p2 requires a correct MAC only if the MAC field has a nonzero length, which makes it easier for man-in-the-middle attackers to spoof packets by omitting the MAC.	3.4.3000
CVE-2015-1799	N/A	The symmetric-key feature in the receive function in ntp_proto.c in ntpd in NTP 3.x and 4.x before 4.2.8p2 performs state-variable updates upon receiving certain invalid packets, which makes it easier for man-in-the-middle attackers to cause a denial of service (synchronization loss) by spoofing the source IP address of a peer.	3.4.3000
CVE-2015-2808	Not vulnerable: Requires strict mode	The RC4 algorithm, as used in the TLS protocol and SSL protocol, does not properly combine state data with key data during the initialization phase, which makes it easier for remote attackers to conduct plaintext-recovery attacks against the initial bytes of a stream by sniffing network traffic that occasionally relies on keys affected by the Invariance Weakness, and then using a brute-force approach involving LSB values, aka the "Bar Mitzvah" issue.	3.4.1120

Table 26 - Common Vulnerabilities and Exposures

CVE	Vulnerability ¹	Description	Fixed in Version
CVE-2015-3456	Not vulnerable	The Floppy Disk Controller (FDC) in QEMU, as used in Xen 4.5.x and earlier and KVM, allows local guest users to cause a denial of service (out-of-bounds write and guest crash) or possibly execute arbitrary code via the (1) FD_CMD_READ_ID, (2) FD_CMD_DRIVE_SPECIFICATION_COMMAND, or other unspecified commands, aka VENOM. Though the VENOM vulnerability is also agnostic of the guest operating system, an attacker (or an attacker's malware) would need to have administrative or root privileges in the guest operating system in order to exploit VENOM.	3.4.3000
CVE-2015-4000	Not vulnerable	The TLS protocol 1.2 and earlier, when a DHE_EXPORT ciphersuite is enabled on a server but not on a client, does not properly convey a DHE_EXPORT choice, which allows man-in-the-middle attackers to conduct cipher-downgrade attacks by rewriting a ClientHello with DHE replaced by DHE_EXPORT and then rewriting a ServerHello with DHE_EXPORT replaced by DHE, aka the "Logjam" issue.	3.4.3000
CVE-2015-5119	Not vulnerable	This candidate has been reserved by an organization or individual that will use it when announcing a new security problem. When the candidate has been publicized, the details for this candidate will be provided.	3.4.3000

1. Vulnerability may take the following three values:

N/A – not relevant to MLNX-OS

Not vulnerable – Mellanox products are protected against this vulnerability

Not vulnerable: Requires strict mode – working in strict mode protects against this vulnerability

Appendix C: UI Changes in Version 3.4.2008



Relevant changes in the CLI in this section are marked in boldface.

In order to improve user interface and unify look and feel across all Mellanox switch platforms, MLNX-OS® versions 3.4.2008 and above introduce some changes in user interface.

C.1 Interface Addressing Change

Interface addressing in interface specific commands has changed in order to align the InfiniBand interface schema with that of Ethernet's, and to support the EDR director switch system which has 2 ASICs per leaf.

C.1.1 CLI Change

Interface referencing for director switches has become <slot/module/port> and their show command output text displays IB<slot/module/port>.

<u>Before:</u>	<u>After:</u>
switch (config) # show interfaces ib L01/1	switch (config) # show interfaces ib 1/1/1
Slot L01 port 1 state	IB1/1/1 state:
Logical port state : Initialize	Logical port state : Initialize
Physical port state : LinkUp	Physical port state : LinkUp
Current line rate : 40.0 Gbps	Current line rate : 40.0 Gbps
Supported speeds : 2.5, 5.0, 10.0(QDR),	Supported speeds : sdr, ddr, qdr, fdr10
10.0(FDR10) or 14.0 Gbos rate	...
...	

Internal interface referencing for director switch systems has become as follows:

<u>Before:</u>	<u>After:</u>
switch (config) # show interfaces ib internal S01/1	switch (config) # show interfaces ib internal spine 1/1/1
Slot S01 port 1 state	IB1/1/1 state:
Connected to slot/chip : 1/1	Connected to slot/chip : 1/1
Connected to port : 19	Connected to port : 19
Connected device active: -	Connected device active: -
Error state : 0	Error state : 0
Logical port state : Initialize	Logical port state : Initialize
Physical port state : LinkUp	Physical port state : LinkUp
Current line rate : 40.0 Gbps	Current line rate : 40.0 Gbps
Supported speeds : .5, 5.0, 10.0(QDR),	Supported speeds : sdr, ddr, qdr, fdr10
10.0(FDR10) or 14.0 Gbos rate	...
...	

Interface show command output text on 1U switches has become IB<slot/port>.

Before:	After:
switch (config) # show interfaces ib 1/1	switch (config) # show interfaces ib 1/1
Slot 1 port 1 state	IB1/1 state:
Logical port state : Down	Logical port state : Down
Physical port state : Polling	Physical port state : Polling
Current line rate : -	Current line rate : -
Supported speeds : 2.5, 5.0, 10.0(QDR),	Supported speeds : sdr, ddr, qdr, fdr10,
10.0(FDR10) or 14.0 Gbos rate	fdr
...	...

C.1.2 MIB ifTable Change

The ifDescr column now displays interfaces in the syntax IB<slot/port>.

Figure 32: 1U MIB ifTable Before Screenshot

	ifIndex	ifDescr
1	1	lo
2	2	mgmt0
3	3	mgmt1
4	4	ib0
5	65	/SX/9
6	67	/SX/12
7	69	/SX/11
8	70	/SX/14
9	72	/SX/13
10	74	/SX/16
11	75	/SX/15
12	77	/SX/18
13	79	/SX/17
14	80	/SX/28

Figure 33: 1U MIB ifTable After Screenshot

	ifIndex	ifDescr
1	1	lo
2	2	mgmt0
3	3	mgmt1
4	4	ib0
5	65	IB1/5
6	69	IB1/6
7	70	IB1/7
8	75	IB1/8
9	79	IB1/9
10	80	IB1/13
11	85	IB1/12
12	89	IB1/11
13	91	IB1/10
14	97	IB1/4

For director switches the ifDescr column displays interfaces in the syntax IB<slot/module/port>.

Figure 34: Director Switch MIB ifTable Before Screenshot

	ifIndex	ifDescr
1	1	lo
2	2	mgmt0
3	4	ib0
4	1616	/L01/10
5	1618	/L01/7
6	1620	/L01/8
7	1621	/L01/5
8	1623	/L01/6
9	1625	/L01/3
10	1627	/L01/4
11	1629	/L01/1
12	1631	/L01/2
13	1648	/L01/9
14	1650	/L01/12

Figure 35: Director Switch MIB ifTable After Screenshot

	ifIndex	ifDescr
331	2389	IB1/1/5
332	2391	IB1/1/6
333	2393	IB1/1/3
334	2395	IB1/1/4
335	2397	IB1/1/1
336	2399	IB1/1/2
337	2416	IB1/1/9
338	2418	IB1/1/12
339	2420	IB1/1/11
340	2421	IB1/1/14
341	2423	IB1/1/13
342	2425	IB1/1/16
343	2427	IB1/1/15
344	2429	IB1/1/18

C.1.3 WebUI Ports Page Change

In the “Ports” page of the MLNX-OS® WebUI, the “Port number” field has been modified to reflect the change in the CLI. “Port number” now displays interfaces in the following syntax for director switches: <slot>/<module>/<port>.

C.2 Interface Speed Configuration Change

Interface speed configuration has changed in order to improve user experience when configuring InfiniBand speeds and to support additional permutations for setting allowed speeds.

C.2.1 CLI Change

Interface speed configuration commands accept any of the following speed name combinations: SDR; DDR; QDR; FDR10; FDR; EDR.



Commands with the old syntax, however, are still supported.

<u>Before:</u>	<u>After:</u>
switch (config interface ib 1/1) # speed ?	switch (config interface ib 1/1) # speed ?
1 2.5 Gbps	sdr 10.0 Gbps rate on 4 lane width
3 2.5 or 5.0 Gbps	ddr 20.0 Gbps rate on 4 lane width
5 2.5 or 10.0(QDR) Gbps	qdr 40.0 Gbps rate on 4 lane width
7 2.5, 5.0 or 10.0(QDR) Gbps	fdr10 40.0 Gbps rate on 4 lane width
8 10.0(FDR10) Gbps	fdr 56.0 Gbps rate on 4 lane width
13 2.5, 10.0(QDR) or 10(FDR10) Gbps	force Force a speed vector without the
15 2.5, 5.0, 10.0(QDR) or 10.0(FDR10) Gbps	sdr bit

<u>Before:</u>	<u>After:</u>
switch (config interface ib 1/1) # speed ?	switch (config interface ib 1/1) # speed ?
10 10.0 Gbps rate on 4 lane width	sdr 10.0 Gbps rate on 4 lane width
100 20.0 Gbps rate on 4 lane width	ddr 20.0 Gbps rate on 4 lane width
1000 40.0 Gbps rate on 4 lane width	qdr 40.0 Gbps rate on 4 lane width
10000 40.0 Gbps rate on 4 lane width	fdr10 40.0 Gbps rate on 4 lane width
auto 56.0 Gbps rate on 4 lane width	fdr 56.0 Gbps rate on 4 lane width
	force Force a speed vector without the
	sdr bit

```
switch (config interface ib 1/1) # speed
ddr   fdr   fdr10   force   qdr   sdr
switch (config interface ib 1/1) # speed sdr fdr qdr
switch (config interface ib 1/1) #
```


Show interface commands display speed names instead of the speed figures.

<u>Before:</u>	<u>After:</u>
switch (config) # show interfaces ib L01/1	switch (config) # show interfaces ib 1/1
Slot L01 port 1 state:	IB1/1 state:
Logical port state : Down	Logical port state : Down
Physical port state : Polling	Physical port state : Polling
Current line rate : -	Current line rate : -
Supported speeds : 2.5, 5.0, 10.0(QDR),	Supported speeds : sdr, ddr, qdr, fdr10,
10.0(FDR10) or 14.0 Gbps	fdr
...	...

C.2.2 WebUI Change

In the Ports page of the MLNX-OS WebUI, the “Supported speeds” and “Speed” fields have been modified to reflect the change in the CLI. “Supported speeds” and “Speed” now display speed names instead of the speed figures.

Figure 36: Ports WebUI Page



C.3 IB SM Link Speed Change

IB SM link speed has changed to improve user experience when configuring InfiniBand interface speeds, to support additional permutations for setting allowed speeds, and to display clear default values and negotiation outputs.

IB SM link speed commands accept any of the following speed name combinations: SDR; DDR; QDR; FDR10; FDR; EDR.



Commands with the old syntax, however, are still supported.

Before:

```
switch (config) # ib sm force-link-speed ?
0 Do not modify switch assigned default value
1 Negotiate only 2.5 Gbps rate
3 Negotiate 2.5 or 5.0 Gbps rate
5 Negotiate 2.5 or 10.0 Gbps rate
7 Negotiate 2.5, 5.0, or 10.0 Gbps rate
13 Negotiate 2.5, 10.0(QDR) or 10.0(FDR19) Gbps
15 Negotiate 2.5, 5.0, 10.0(QDR) or 10.0(FDR10) Gbps
21 Negotiate 2.5, 10.0(QDR) or 14.0 Gbps
23 Negotiate 2.5, 5.0, 10.0(QDR) or 14.0 Gbps
29 Negotiate 2.5, 10.0(QDR), 10.0(FDR10) or 14.0 Gbps
31 Negotiate 2.5, 5.0, 10.0(QDR), 10.0(FDR10) or 14.0 Gbps
```

After:

```
switch (config) # ib sm force-link-speed ?
sdr 10.0 Gbps rate on 4 lane width
ddr 20.0 Gbps rate on 4 lane width
qdr 40.0 Gbps rate on 4 lane width
fdr10 40.0 Gbps rate on 4 lane width
fdr 56.0 Gbps rate on 4 lane width
edr 100.0 Gbps rate on 4 lane width
```

The command “no ib sm force-link-speed” configures default speed.

Before:	After:
switch (config) # no ib sm force-link-speed	switch (config) # no ib sm force-link-speed
switch (config) # show ib sm force-link-speed	switch (config) # show ib sm force-link-speed
15 (Negotiate 2.5, 5.0, 10.0(QDR) or 10.0(FDR10) Gbps)	Default: set to PortInfo:LinkSpeedSupported

Show IB SM speed command displays negotiation as well as speed names.

Before:	After:
switch (config) # show ib sm force-link-speed 5 (Negotiate 2.5 or 10.0 Gbps rate)	switch (config) # show ib sm force-link-speed Negotiate: sdr, qdr

The output of the command “show ib sm force-link-speed-ext”:

Before:	After:
<pre>switch (config) # show ib sm force-link-speed-ext 1 (Allow extended negotiation speeds with 14.0 Gbps rate)</pre>	<pre>switch (config) # show ib sm force-link-speed-ext Negotiate: fdr</pre>

C.4 Multi-ASIC Support

Multi-ASIC support has been added to MLNX-OS in order to support the new EDR director switch systems whose leafs feature two ASICs per leaf, to improve MIB module indexing and better represent module hierarchy, to add an additional hierarchy level with an ASIC device to support more than one ASIC per module, and to add support to all sensors.

C.4.1 CLI Change

Added a new “Device” column to the output of the command “show guidb”.

<u>Before:</u>		<u>After:</u>	
switch (config) # show guid		switch (config) # show guid	
=====		=====	
SX module	GUID	Module	Device GUID
=====		=====	
SYSTEM	00:05:C9:03:00:42:D8:00	SYSTEM	- 00:05:C9:03:00:42:D8:00
S01	00:02:C9:03:00:84:3B:60	S01	SX 00:02:C9:03:00:84:3B:60
L17	00:02:C9:03:00:84:3B:40	S02	SX 00:02:C9:03:00:84:3B:40
L18	00:02:C9:03:00:84:3A:F0	S03	SX 00:02:C9:03:00:84:3A:F0
S05	00:02:C9:03:00:84:3B:70	S04	SX 00:02:C9:03:00:84:3B:70
L01	00:02:C9:03:00:84:3B:80	S05	SX 00:02:C9:03:00:84:3B:80
L03	00:02:C9:03:00:84:3B:90	S06	SX 00:02:C9:03:00:84:3B:90
L04	00:02:C9:03:00:61:ED:00	S07	SX 00:02:C9:03:00:61:ED:00
L06	00:02:C9:03:00:84:3B:A0	S08	SX 00:02:C9:03:00:84:3B:A0
L09	00:02:C9:03:00:84:3B:50	S09	SX 00:02:C9:03:00:84:3B:50
L11	00:02:C9:03:00:84:3B:10	S10	SX 00:02:C9:03:00:84:3B:10
L13	00:02:C9:03:00:61:EC:B0	S11	SX 00:02:C9:03:00:61:EC:B0
L15	00:02:C9:03:00:66:C9:B0	S12	SX 00:02:C9:03:00:66:C9:B0
L16	00:02:C9:03:00:66:C9:60	S13	SX 00:02:C9:03:00:66:C9:60
...		...	

The output of the command “show guid” on EDR director switch:

```
switch (config) # show guids
=====
Module      Device      GUID
=====
SYSTEM      -           00:05:C9:03:00:42:D8:00
S01         SIB         00:02:C9:03:00:84:3B:60
S02         SIB         00:02:C9:03:00:84:3B:40
S03         SIB         00:02:C9:03:00:84:3A:F0
S04         SIB         00:02:C9:03:00:84:3B:70
S05         SIB         00:02:C9:03:00:84:3B:80
S06         SIB         00:02:C9:03:00:84:3B:90
S07         SIB         00:02:C9:03:00:61:ED:00
S08         SIB         00:02:C9:03:00:84:3B:A0
S09         SIB         00:02:C9:03:00:84:3B:50
S10         SIB         00:02:C9:03:00:84:3B:10
S11         SIB         00:02:C9:03:00:61:EC:B0
S12         SIB         00:02:C9:03:00:66:C9:B0
S13         SIB         00:02:C9:03:00:66:C9:60
S14         SIB         00:02:C9:03:00:66:C9:50
S15         SIB         00:02:C9:03:00:66:C9:50
S16         SIB         00:02:C9:03:00:66:C9:50
101         SIB1        00:02:C9:03:00:31:81:80
101         SIB2        00:02:C9:03:00:31:81:81
...
```

On 1U switches, the command “show guids” displays “MGMT” under the “Module” column instead of “1”.

<u>Before:</u>		<u>After:</u>	
switch (config) # show guids		switch (config) # show guids	
=====		=====	
SX module	GUID	Module	Device GUID
=====		=====	
SYSTEM	00:02:C9:03:00:A8:EA:10	SYSTEM	- F4:52:14:03:00:11:E4:F0
1	00:02:C9:03:00:A8:EA:12	MGMT	SX F4:52:14:03:00:11:E4:F2

Added a “Device” column to the command “show asic-version”.

<u>Before:</u>		<u>After:</u>	
switch (config) # show asic-version		switch (config) # show asic-version	
=====		=====	
Module	Version	Module	Device Version
=====		=====	
MGMT	9.3.3150	MGMT	SX 9.3.3150

The output of the command “show asic-version” on an EDR director switch system:

```
switch (config) # show guids
=====
Module           Device           Version
=====
101              SIB1             11.0.1296
101              SIB2             11.0.1296
102              SIB1             11.0.1296
102              SIB2             11.0.1296
103              SIB1             11.0.1296
103              SIB2             11.0.1296
104              SIB1             11.0.1296
104              SIB2             11.0.1296
105              SIB1             11.0.1296
105              SIB2             11.0.1296
106              SIB1             11.0.1296
106              SIB2             11.0.1296
...
```

C.4.2 MIB entPhysicalTable Change

The entPhysicalIndex and entPhysicalDescr columns now display and convey module hierarchy.

- entPhysicalDescr in entPhysicalTable revamped now represents ASIC module hierarchy
- entPhysicalIndex in entPhysicalTable now represents a legend and not just a running number. For example (line 9 in [Figure 38](#)), “S01/BOARD_MONITOR/T1” has the index 301030011 which indicates the following: 3=Spine, 1=index, 3=BOARD_MONITOR, 1=T, 1=T1.

Figure 37: MIB entPhysicalTable Before Screenshot

	entPhysicalIndex	entPhysicalDescr
1	1	Mellanox SX6512, 216-Port FDR/FDR10 Switch System
2	2	MGMT1
3	3	MGMT2
4	4	S01
5	5	S02
6	6	S03
7	7	S04
8	8	S05
9	9	S06
10	10	L01
11	11	L02
12	12	L03
13	13	L04
14	14	L05

Figure 38: MIB entPhysicalTable After Screenshot

	entPhysicalIndex	entPhysicalDescr
1	1	Mellanox SX6518, 324-Port FDR/FDR10 Switch System
2	201000000	MGMT1
3	201040011	MGMT1/CPU_BOARD_MONITOR/T1
4	201040012	MGMT1/CPU_BOARD_MONITOR/T2
5	202000000	MGMT2
6	301000000	S01
7	301010021	S01/FAN/F1
8	301010022	S01/FAN/F2
9	301030011	S01/BOARD_MONITOR/T1
10	301050000	S01/SX
11	301050011	S01/SX/T1
12	301100011	S01/SX_AMBIENT_TEMP/T1
13	302000000	S02
14	302010021	S02/FAN/F1

C.5 MGMT Module Display Change

The MGMT module display is improved to better represent the actual structure of modules within the system. The commands “show power consumers” and “show temperature” now display information from the MGMT module.

The output of the command “show power consumers” on a 1U PPC switch system:

Before:						After:					
switch (config) # show power consumers						switch (config) # show power consumers					
=====						=====					
Module	Device	Power (Watts)	Voltage	Current (Amp)	Status	Module	Device	Power (Watts)	Voltage	Current (Amp)	Status
=====						=====					
CURR_MONITOR	MONITOR	33.31	11.72	2.84	OK	MGMT	CURR	33.31	11.72	2.84	OK
Total power used : 33.31 W						Total power used : 33.31 W					
Max power : 235.00 W						Max power : 235.00 W					

The output of the command “show temperature” on a 1U PPC switch system:

Before:					After:				
switch (config) # show temperature					switch (config) # show temperature				
=====					=====				
Module	Component	Reg	CurTemp (Celsius)	Status	Module	Component	Reg	CurTemp (Celsius)	Status
=====					=====				
MGMT	BOARD MONITOR	T1	26.50	OK	MGMT	CPU Core Sensor	T1	28.00	OK
MGMT	CPU_MEZZ_TEMP	T1	27.00	OK	MGMT	CPU Core Sensor	T2	30.00	OK
CPU_X86	CPU Core Sensor	T1	28.00	OK	MGMT	CPU Core Sensor	T3	56.00	OK
CPU_X86	CPU Core Sensor	T2	30.00	OK	MGMT	CPU Core Sensor	T4	26.00	OK
CPU_X86	CPU Core Sensor	T3	56.00	OK	MGMT	CPU package Sensor	T4	40.00	OK
CPU_X86	CPU Core Sensor	T4	26.00	OK	MGMT	BOARD_MONITOR	T1	26.50	OK
CPU_X86	CPU package Sensor	T4	38.00	OK	MGMT	CPU_MEZZ_TEMP	T1	27.00	OK
MGMT	QSFP_TEMP1	T1	27.00	OK	MGMT	QSFP_TEMP1	T1	27.00	OK
...					...				

C.6 MLNX-OS Image Name Change

The “SX_” prefix has been removed from the name of the MLNX-OS image.

C.6.1 CLI Change

The output of the command “show version”:

<u>Before:</u>	<u>After:</u>
switch (config) # show version	switch (config) # show version
Product name: SX_X86_64	Product name: MLNX-OS
Product release: SX_3.4.0008	Product release: 3.4.2000
Build ID: #1-dev	Build ID: #1-dev
Build date: 2014-11-10 20:07:51	Build date: 2015-05-06 02:16:39
Target arch: x86_64	Target arch: x86_64
Target hw: x86_64	Target hw: m460ex
Built by: jenkins@fit74	Built by: jenkins@fit74
Version summary: SX_X86_64 SX_3.4.0008 2014-11-10 20:07:51 X86_64	Version summary: X86_64 3.4.2000 2015-04-12 20:06:05 X86_64
...	...

The output of the command “show version”:

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

Partition 1:
PPC_M460EX 3.4.2000 2015-05-06 02:16:39 ppc

Partition 2:
SX_PPC_M460EX SX_3.4.0000 2014-10-14 20:26:41 ppc //older version with "SX_" suffix

Last boot partition: 1
Next boot partition: 1
...
```

C.6.2 WebUI Status Page Change

In the “Status” page of the MLNX-OS WebUI, the “Software Version” field has been modified to reflect the change in the image file name so now it appears without the prefix “SX_” ([Figure 39](#)).

Figure 39: Status WebUI Page

The screenshot displays the Status WebUI Page. On the left is a navigation menu with links: Summary, Profile and Capabilities, Temperature, Power Supplies, Fans, CPU Load, Memory, Network, Logs, Maintenance, Alerts, and Virtualization. The main content area is titled 'Summary' and contains a table of system information. A red arrow points to the 'Software Version' field. Below the summary table is a section for 'Active alerts' which currently shows 'No alerts'. At the bottom of the page is a copyright notice: '© 2009-2015 Mellanox Technologies, Inc.'

Summary	
Date and Time:	2015/04/15 11:52:17
Hostname:	Piranha-157-x86
Uptime:	8h 23m 54s
Software Version:	X86_64 3.4.1890_M 2015-04-14 22:04:05 x86_64
Model:	x86
Host ID:	F45214C962BE
System memory:	1666 MB used / 6258 MB free / 7924 MB total
CPU load averages:	1.00 / 1.01 / 1.05
System UUID	03000200-0400-0500-0006-000700080009

Active alerts

No alerts

© 2009-2015 Mellanox Technologies, Inc.

C.7 CPU Module Display Change

The CPU module has been removed from the outputs of the CLI commands “show inventory” and “show modules”.

C.7.1 CLI Change

The command “show inventory” does not list CPU under the “Module” column.

Before

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

Module	Type	Part number	Serial Number	Asic revision	HW Revision
CHASSIS	SX6036	MSX6036F-1SFS	MT1343X02004	N/A	A3
MGMT	SX6036	MSX6036F-1SFS	MT1343X02004	2	A3
FAN	SXX0XX_FAN	MSX60-FF	MT1342X03954	N/A	A1
PS1	SXX0XX_PS	MSX60-PF	MT1342X03824	N/A	A1
CPU	CPU	SA000203-B	MT1140X00201	N/A	A1

After

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

Module	Type	Part number	Serial Number	Asic revision	HW Revision
CHASSIS	SX6036	MSX6036F-1SFS	MT1343X02004	N/A	A3
MGMT	SX6036	MSX6036F-1SFS	MT1343X02004	2	A3
FAN	SXX0XX_FAN	MSX60-FF	MT1342X03954	N/A	A1
PS1	SXX0XX_PS	MSX60-PF	MT1342X03824	N/A	A1

The command “show module” does not list CPU under the “Module” column.

Before:					After:				
switch (config) # show module					switch (config) # show module				
Module	Type	Present	Power	Is Fatal	Module	Type	Present	Power	Is Fatal
MGMT	SX6036	1	1	Not Fatal	MGMT	SX6036	1	1	Not Fatal
FAN	SXX0XX_FAN	1	1	Not Fatal	FAN	SXX0XX_FAN	1	1	Not Fatal
PS1	SXX0XX_PS	1	1	Not Fatal	PS1	SXX0XX_PS	1	1	Not Fatal
PS2	SXX0XX_PS	0	0	Not Fatal	PS2	SXX0XX_PS	0	0	Not Fatal
CPU	CPU	1	1	Not Fatal					

C.7.2 WebUI System Inventory Page Change

In the “System > Inventory” page of the MLNX-OS WebUI, the “Module” column does not display the CPU module anymore (Figure 40).

Figure 40: System Inventory WebUI Page

Modules		Modules Information						
Inventory		Module	Type	Part number	Serial number	Asic revision	HW revision	Node GUID
Power Management								Asic FW version
MLNX-OS Upgrade		CHASSIS	MSX1700	MSX1700-BS1F2	MT1341X00736	N/A	A3	F4:52:14:03:00:0F:50:B0 -
Reboot		MGMT	MSX1700	MSX1700-BS1F2	MT1341X00736	2	A3	- 9.3.3090
		FAN1	SX1024_FAN	MSX62-FF	MT1344X00951	N/A	A1	-
		FAN2	SX1024_FAN	MSX62-FF	MT1344X00944	N/A	A1	-
		PS2	400W_PS	MSX64-PF	MT1331X00028	N/A	A1	-
		CPLDs Information						
		Name	Type	Version				
		Cpld1	CPLD_TOR	5				
		Cpld2	CPLD_PORT1	3				
		Cpld3	CPLD_MEZZ	4				

© 2009-2015 Mellanox Technologies, Inc.