



NVIDIA MLNX-OS User Manual

v3.11.2224 LTS

Table of Contents

1	Overview	11
1.1	Intended Audience	11
1.2	Related Documentation	11
1.3	Terminology	11
1.4	System Features.....	13
1.5	InfiniBand Features	14
2	Getting Started.....	15
2.1	Configuring the Switch for the First Time.....	15
2.1.1	Configuring the Switch with ZTP	22
2.1.2	Rerunning the Wizard	22
2.2	Starting the Command Line (CLI).....	22
2.3	Starting the Web User Interface (WebUI)	23
2.4	Zero-touch Provisioning	25
2.4.1	Running DHCP-ZTP	25
2.4.2	ZTP on Modular Switches	27
2.4.3	ZTP and OS Upgrade	27
2.4.4	DHCPv4 Configuration Example.....	28
2.4.5	DHCPv6 Configuration Example.....	28
2.4.6	ZTP Commands	28
2.5	Licenses	30
2.5.1	Installing OS License via CLI.....	30
2.5.2	Installing OS License via Web	30
2.5.3	Retrieving a Lost License Key	32
2.5.4	Additional Reading and Use Cases.....	32
2.5.5	License Commands	33
3	User Interfaces	35
3.1	LED Indicators	35
3.2	Command Line Interface (CLI).....	35
3.2.1	CLI Modes	35
3.2.2	Syntax Conventions.....	36
3.2.3	Getting Help	36
3.2.4	Prompt and Response Conventions.....	37

3.2.5	Using the “no” Command Form.....	38
3.2.6	Parameter Key.....	38
3.2.7	CLI Pipeline Operator Commands	39
3.3	Secure Shell (SSH)	42
3.3.1	Adding a Host and Providing an SSH Key	42
3.3.2	Retrieving Return Codes When Executing Remote Commands	43
3.4	Web Interface Overview	43
3.4.1	Password Hardening	43
3.4.2	Changing Default Password.....	44
3.4.3	About Web UI	45
3.4.4	Setup Menu	46
3.4.5	System Menu.....	47
3.4.6	Security Menu	47
3.4.7	Ports Menu	48
3.4.8	Status Menu	48
3.4.9	IB SM Mgmt Menu	49
3.4.10	IB Router Menu	50
3.5	UI Commands	50
3.5.1	CLI Session	50
3.5.2	Web Interface	69
4	System Management.....	77
4.1	Management Interfaces	77
4.1.1	Configuring Management Interfaces with Static IP Addresses.....	77
4.1.2	Configuring IPv6 Address on the Management Interface.....	77
4.1.3	Dynamic Host Configuration Protocol (DHCP)	78
4.1.4	Default Gateway	78
4.1.5	Configuring Hostname via DHCP (DHCP Client Option 12).....	78
4.1.6	Management Interface Commands	79
4.1.7	Control Plane Policing (CoPP).....	100
4.2	Chassis Management.....	113
4.2.1	System Health Monitor	113
4.2.2	Power Management.....	115
4.2.3	Monitoring Environmental Conditions.....	117
4.2.4	USB Access	118

4.2.5	Unit Identification LED.....	119
4.2.6	High Availability (HA).....	119
4.2.7	System Reboot.....	122
4.2.8	Viewing Active Events	122
4.2.9	Chassis Management Commands.....	123
4.3	UNBREAKABLE-LINK® Adapter and Switch Technology	140
4.3.1	LLR Mode	141
4.3.2	LLR Negotiation	141
4.3.3	LLR Status	141
4.3.4	UNBREAKABLE-LINK® Switch Commands	141
4.4	Upgrade/Downgrade Process.....	142
4.4.1	Important Pre-OS Upgrade Notes	142
4.4.2	Upgrading Operating System Software	143
4.4.3	Upgrading HA Groups	145
4.4.4	Upgrading MLNX-OS Software on Modular Switches	146
4.4.5	Deleting Unused Images	147
4.4.6	Downgrading OS Software	147
4.4.7	Upgrading System Firmware	150
4.4.8	Software Management Commands	152
4.5	Configuration Management.....	158
4.5.1	Saving a Configuration File.....	158
4.5.2	Loading a Configuration File	158
4.5.3	Restoring Factory Default Configuration	159
4.5.4	Managing Configuration Files	159
4.5.5	Automated Periodic Configuration File Backup	161
4.5.6	Configuration Management Commands.....	162
4.6	mDNS.....	180
4.6.1	mDNS Commands.....	180
5	NTP and Clock	181
5.1	NTP Authenticate	181
5.2	NTP Authentication Key	181
5.3	NTP Commands.....	181
5.3.1	clock set.....	182
5.3.2	clock timezone	183

5.3.3	ntp	183
5.3.4	ntpdate	184
5.3.5	ntp authenticate	184
5.3.6	ntp authentication-key	184
5.3.7	ntp peer disable.....	185
5.3.8	ntp peer keyID.....	185
5.3.9	ntp peer version.....	186
5.3.10	ntp server disable.....	186
5.3.11	ntp server keyID.....	187
5.3.12	ntp server-role disable	187
5.3.13	ntp server trusted-enable	187
5.3.14	ntp server version.....	188
5.3.15	ntp trusted-key.....	188
5.3.16	show clock	189
5.3.17	show ntp.....	189
5.3.18	show ntp configured.....	189
5.3.19	show ntp keys	190
6	Network Management Interfaces	191
6.1	SNMP	191
6.1.1	Standard MIBs.....	191
6.1.2	Private MIBs	192
6.1.3	Proprietary Traps.....	192
6.1.4	Configuring SNMP	193
6.1.5	Resetting SNMPv3 Engine ID	193
6.1.6	Configuring an SNMPv3 User.....	194
6.1.7	Configuring SNMP Notifications (Traps or Informs)	195
6.1.8	SNMP SET Operations.....	196
6.1.9	Additional Readings and Use Cases.....	200
6.2	JSON API.....	200
6.2.1	Authentication	200
6.2.2	Sending the Request.....	203
6.2.3	JSON Request Format	203
6.2.4	JSON Response Format	205
6.2.5	Supported Commands	207

6.2.6	JSON Examples	207
6.2.7	JSON Request Using WebUI	211
6.2.8	Additional Reading and Use Cases.....	213
6.3	Network Management Interface Commands	213
6.3.1	SNMP	214
6.3.2	JSON API.....	225
7	Virtualization	227
7.1	Limiting the Container's Resources	227
7.1.1	Memory Resources Allocation Protocol	228
7.1.2	CPU Resource Allocation Protocol	228
7.2	Upgrade Ramifications	228
7.2.1	Changing Docker Storage Driver	228
7.3	Docker Containers Commands	229
7.3.1	docker	229
7.3.2	docker login.....	230
7.3.3	docker logout.....	230
7.3.4	commit	230
7.3.5	copy-sdk.....	231
7.3.6	remove image	231
7.3.7	exec	232
7.3.8	label	232
7.3.9	load	232
7.3.10	pull.....	233
7.3.11	save	233
7.3.12	shutdown	234
7.3.13	start.....	234
7.3.14	image upload	235
7.3.15	file image upload	236
7.3.16	show docker	236
7.3.17	show docker containers.....	236
7.3.18	show docker images	238
7.3.19	show docker ps.....	238
7.3.20	show docker labels	239
7.3.21	show docker login.....	239

7.3.22	show docker stats.....	239
8	Telemetry, Monitoring, and Debuggability	241
8.1	Logging	241
8.1.1	Monitor	241
8.1.2	Remote Logging	241
8.1.3	Logging Protocol	242
8.1.4	Logging Commands	242
8.2	Link Diagnostic Per Port.....	257
8.2.1	Link Diagnostic Commands	259
8.3	Event Notifications	260
8.3.1	Supported Event Notifications and MIB Mapping.....	260
8.3.2	SNMP Trap Notification.....	262
8.3.3	Terminal Notifications.....	262
8.3.4	Email Notifications	263
8.3.5	Command Event Notifications	263
8.4	Buffer Histograms Monitoring	274
8.4.1	Buffer Histograms and Thresholds Commands.....	275
8.5	Statistics and Alarms	285
8.5.1	Commands	285
8.6	Management Information Bases (MIBs)	297
8.6.1	Calculating of entPhysicalIndex in the Entity MIB	297
8.6.2	Examples	299
9	User Management, Authentication, & Security.....	300
9.1	User Management & Security	300
9.1.1	User Accounts.....	300
9.1.2	Authentication, Authorization, and Accounting (AAA)	300
9.1.3	User Re-authentication	301
9.1.4	RADIUS.....	301
9.1.5	TACACS+	301
9.1.6	LDAP	301
9.1.7	System Secure Mode.....	302
9.1.8	User Management and Security Commands	304
9.2	Cryptographic (X.509, IPSec) and Encryption	331
9.2.1	System File Encryption.....	331

9.2.2	Changing a Certificate in the System	331
9.2.3	Cryptographic and Encryption Commands	333
10	InfiniBand Switching.....	342
10.1	Node Name.....	342
10.1.1	Node Name Commands.....	342
10.2	Fabric	343
10.2.1	Fabric Commands	343
10.3	IB Router	344
10.3.1	Configuring IB Router.....	346
10.3.2	Subnet Prefix Checking	347
10.3.3	IB Router Commands	348
10.4	InfiniBand Interface	351
10.4.1	Transceiver Information.....	351
10.4.2	High Power Transceivers	351
10.4.3	Forward Error Correction	351
10.4.4	Breakout Cables.....	352
10.4.5	InfiniBand Interface Commands	354
10.5	Subnet Manager	366
10.5.1	Partitions	367
10.5.2	Adaptive Routing.....	368
10.5.3	Scatter Ports.....	368
10.5.4	GUID Routing Order.....	368
10.5.5	Bulk Update Mode	369
10.5.6	SM Commands	369
10.6	Subnet Manager High Availability.....	443
10.6.1	Joining, Creating or Leaving an InfiniBand Subnet ID	444
10.6.2	MLNX-OS Management Centralized Location	445
10.6.3	High Availability Node Roles.....	445
10.6.4	Configuring MLNX-OS SM HA Centralized Location	446
10.6.5	Creating and Adding Systems to an InfiniBand Subnet ID	446
10.6.6	Restoring Subnet Manager Configuration.....	446
10.6.7	SM HA Commands	448
11	Appendixes.....	452
11.1	Appendix: Enhancing System Security According to NIST SP 800-131A	452

11.1.1	Web Certificate	452
11.1.2	SNMP	453
11.1.3	HTTPS.....	454
11.1.4	Code Signing	455
11.1.5	SSH	455
11.1.6	LDAP	456
11.2	Appendix: Splunk Integration with NVIDIA Products.....	457
11.2.1	Getting Started with Splunk.....	457
11.2.2	Switch Configuration	457
11.2.3	Adding a Task	458
11.2.4	Retrieving Data from TCP and UDP Ports.....	459
11.2.5	SNMP Input to Poll Attribute Values and Catch Traps.....	461
11.3	Appendix: Show Commands Not Supported By JSON API.....	464
12	Document Revision History	466

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

Welcome to MLNX-OS Documentation

NVIDIA® MLNX-OS® operating system, enables the management and configuration of NVIDIA's InfiniBand switch system platforms.

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

These pages provide information about the scope, organization, and command line interface of MLNX-OS as well as configuration examples.

Software Download

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

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

Technical Support

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

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

Customers who purchased NVIDIA M-1 Global Support Services, please see your contract for details regarding Technical Support.

Customers who purchased NVIDIA products through an NVIDIA-approved reseller should first seek assistance through their reseller.

Document Revision History

A list of the changes made to the User Manual are provided in [User Manual Revision History](#).

1 Overview



1.1 Intended Audience

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

1.2 Related Documentation

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

Document Name	Description
System Hardware User Manual	This document contains hardware descriptions, LED assignments, and hardware specifications, among other things
Switch Product Release Notes	Please look up the relevant switch system/series release note file

1.3 Terminology


Term	Description
AAA	Authentication, Authorization, and Accounting: <ul style="list-style-type: none">• Authentication—verifies user credentials (username and password)• Authorization—grants or refuses privileges to a user/client for accessing specific services• Accounting—tracks network resources consumption by users
ARP	Address Resolution Protocol. A protocol that translates IP addresses into MAC addresses for communication over a local area network (LAN).
CLI	Command Line Interface. A user interface in which you type commands at the prompt.
DCBX	Domain Name System. A hierarchical naming system for devices in a computer network.
DHCP	The Dynamic Host Configuration Protocol (DHCP) is an automatic configuration protocol used on IP networks.
Modular switch	A high density InfiniBand chassis switch system.
DNS	Domain Name System. A hierarchical naming system for devices in a computer network.
Fabric management	The use of a set of tools (APIs) to configure, discover, and manage and a group of devices organized as a connected fabric.
FTP/TFTP/sFTP	File Transfer Protocol (FTP) is a standard network protocol used to transfer files from one host to another over a TCP-based network, such as the Internet.
Gateway	A network node that interfaces with both InfiniBand and Ethernet, using different network protocols.
GID	Global Identifier. A 128-bit number used to identify a Port on a network adapter (see below), a port on a Router, or a Multicast Group.

Term	Description
GUID	Globally Unique Identifier. A 64-bit number that uniquely identifies a device or component in a subnet.
HA	High Availability. A system design protocol that provides redundancy of system components, thus enables overcoming single or multiple failures in minimal downtime.
Host	A computer platform executing an Operating System which may control one or more network adapters.
IB	InfiniBand
LID	Local Identifier. A 16 bit address assigned to end nodes by the subnet manager. Each LID is unique within its subnet.
LLDP	Link Layer Discovery Protocol. A vendor neutral link layer protocol used by network devices to advertise their identify, capabilities and for neighbor discovery.
MAC	A Media Access Control address (MAC address) is a unique identifier assigned to network interfaces for communications on the physical network segment. MAC addresses are used for numerous network technologies and most IEEE 802 network technologies including Ethernet.
MTU	Maximum Transfer Unit. The maximum size of a packet payload (not including headers) that can be sent /received from a port.
Network Adapter	A hardware device that allows for communication between computers in a network.
RADIUS	Remote Authentication Dial In User Service. A networking protocol that enables AAA centralized management for computers to connect and use a network service.
RDMA	Remote Direct Memory Access. Accessing memory in a remote side without involvement of the remote CPU.
SA	Subnet Administrator (SA) is the interface for querying and manipulating subnet management data.
SCP	Secure Copy or SCP is a means of securely transferring computer files between a local and a remote host or between two remote hosts. It is based on the Secure Shell (SSH) protocol.
SM	Subnet Manager. An entity that configures and manages the subnet, discovers the network topology, assign LIDs, determines the routing schemes and sets the routing tables. There is only one master SM and possible several slaves (Standby mode) at a given time. The SM administers switch routing tables thereby establishing paths through the fabric.
SNMP	Simple Network Management Protocol. A network protocol for the management of a network and the monitoring of network devices and their functions.
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.

1.4 System Features

Feature	Detail
Software management	<ul style="list-style-type: none">• Dual software image• Software and firmware updates• Docker
File management	<ul style="list-style-type: none">• FTP• TFTP• SCP
Logging	<ul style="list-style-type: none">• Event history log• SysLog support
Management interface	<ul style="list-style-type: none">• DHCP/Zeroconf• IPv6
Chassis management	<ul style="list-style-type: none">• Monitoring environmental controls• Power management• Auto-temperature control• High availability
Network management interfaces	<ul style="list-style-type: none">• SNMP v1,v2c,v3• JSON
Security	<ul style="list-style-type: none">• SSH• Telnet• RADIUS• TACACS+
Date and time	<ul style="list-style-type: none">• NTP
Cables & transceivers	<ul style="list-style-type: none">• Transceiver info
Unbreakable links	<ul style="list-style-type: none">• LLR

1.5 InfiniBand Features



Feature	Detail
Subnet manager	<ul style="list-style-type: none"> • OpenSM • Partitions • High availability

2 Getting Started



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

2.1 Configuring the Switch for the First Time

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


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

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

To initialize the switch do the following:

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

The console ports for systems are shown below.



QM87xx Managed System



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

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

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

Using NVIDIA cables is mandatory.

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

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

3. The boot menu is prompted.

```
...
This terminal is not active for input or output while booting.

Boot Menu
-----
0: <image #1>
1: <image #2>
-----

Use the ^ and v keys to select which entry is highlighted.
Press enter to boot the selected image or 'p' to enter a
password to unlock the next set of features.

Highlighted entry is 0:
```


Select "0" to boot with software version installed on partition #1.
 Select "1" to boot with software version installed on partition #2.

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

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

IP configuration by DHCP:

Wizard Session Display (Example)	Comments
Do you want to use the wizard for initial configuration? yes	You must perform this configuration the first time you operate the switch or after resetting the switch to the factory defaults. Type "yes" and then press <Enter>.
Step 1: Hostname? [switch-1]	If you wish to accept the default hostname, then press <Enter>. Otherwise, type a different hostname and press <Enter>.
Step 2: Use DHCP on mgmt0 interface? [yes]	Perform this step to obtain an IP address for the switch. (mgmt0 is the management port of the switch.) - If you wish the DHCP server to assign the IP address, type "yes" and press <Enter>. If you type "no" (no DHCP), then you will be asked whether you wish to use the "zeroconf" configuration or not. If you enter "yes" (yes Zeroconf), the session will continue as shown in the "IP zeroconf configuration" table . If you enter "no" (no Zeroconf), then you need to enter a static IP, and the session will continue as shown in the "Static IP configuration" table .
Step 3: Enable IPv6 [yes]	Perform this step to enable IPv6 on management ports. The default is "yes" (enabled). If you enter "no" (no IPv6), then you will automatically be referred to Step 5.
Step 4: Enable IPv6 autoconfig (SLAAC) on mgmt0 interface? [no]	Perform this step to enable stateless address autoconfig on external management port. The default is "no" (disabled). If you wish to enable it, type "yes" and press <Enter>.
Step 5: Use DHCPv6 on mgmt0 interface? [yes]	Perform this step to enable DHCPv6 on the MGMT0 interface.
Step 6: Update time?	Perform this step to change the time configured. Press <enter> to leave the current time.

Wizard Session Display (Example)	Comments
Step 7: Enable password hardening? [yes]	Perform this step to enable/disable password hardening on your machine. If enabled, new passwords will be checked upon configured restrictions. The default is "yes" (enabled). If you wish to disable it, enter "no".
Step 8: Admin password (Must be typed)? <new_password>	To avoid illegal access to the machine, please type a password and then press <Enter>. Starting from the 3.8.2000 release, the user must type in the admin password upon initial configuration. Due to Senate Bill No. 327, this stage is required and cannot be skipped.
Step 9: Confirm admin password? <new_password>	Confirm the password by re-entering it. Note that password characters are not printed.
Step 10: Monitor password (Must be typed)? <new_password>	To avoid illegal access to the machine, please type a password and then press <Enter>. Starting from the 3.8.2000 release, the user must type in the admin password upon initial configuration. Due to Senate Bill No. 327, this stage is required and cannot be skipped.
Step 11: Confirm monitor password? <new_password>	Confirm the password by re-entering it. Note that password characters are not printed.
You have entered the following information: Hostname: <switch name> Use DHCP on mgmt0 interface: yes Enable IPv6: yes Enable IPv6 autoconfig (SLAAC) on mgmt0 interface: yes Enable DHCPv6 on mgmt0 interface: no Update time: <current time> Enable password hardening: yes Admin password (Enter to leave unchanged): (CHANGED) To change an answer, enter the step number to return to. Otherwise hit <enter> to save changes and exit. Choice: <Enter> Configuration changes saved. To return to the wizard from the CLI, enter the "configuration jump-start" command from configuration mode. Launching CLI... <switch name> [standalone: master] >	The wizard displays a summary of your choices and then asks you to confirm the choices or to re-edit them. Either press <Enter> to save changes and exit, or enter the configuration step number that you wish to return to. To run the command "configuration jump-start" you must be in Config mode.

IP configuration by DHCP for modular switch systems:

Wizard Session Display (Example)	Comments
Do you want to use the wizard for initial configuration? yes	You must perform this configuration the first time you operate the switch or after resetting the switch to the factory defaults. Type "y" and then press <Enter>.
Step 1: Hostname? [switch-1]	If you wish to accept the default hostname, then press <Enter>. Otherwise, type a different hostname and press <Enter>.

Wizard Session Display (Example)	Comments
Step 2: Use DHCP on mgmt0 interface? [yes]	<p>Perform this step to obtain an IP address for the switch. (mgmt0 is the management port of the switch.)</p> <p>If you wish the DHCP server to assign the IP address, type “yes” and press <Enter>.</p> <p>If you type “no” (no DHCP), then you will be asked whether you wish to use the “zeroconf” configuration or not. If you enter “yes” (yes Zeroconf), the session will continue as shown in the IP zeroconf configuration table.</p> <p>If you enter “no” (no Zeroconf), then you need to enter a static IP, and the session will continue as shown in the Static IP configuration table.</p>
Step 3: Enable IPv6 [yes]	<p>Perform this step to enable IPv6 on management ports.</p> <p>If you wish to enable IPv6, type “yes” and press <Enter>.</p> <p>If you enter “no” (no IPv6), then you will automatically be referred to Step 5.</p>
Step 4: Enable IPv6 autoconfig (SLAAC) on mgmt0 interface	<p>Perform this step to enable StateLess address autoconfig on external management port.</p> <p>If you wish to enable it, type “yes” and press <Enter>.</p> <p>If you wish to disable it, enter “no”.</p>
Step 5: Use DHCPv6 on mgmt0 interface? [yes]	<p>Perform this step to enable DHCPv6 on the MGMT0 interface.</p>
Step 6: Admin password (Press <Enter> to leave unchanged)? <new_password>	<p>To avoid illegal access to the machine, please type a password and then press <Enter>.</p>
Step 7: Confirm admin password? <new_password> (this step only happens if you change the password)	<p>Confirm the password by re-entering it.</p> <p>Note that password characters are not printed.</p>
Step 9: HA Chassis Management IP netmask? (Example: [255.255.255.0])	<p>Perform this step to configure the box IPv4 netmask.</p> <p>If you wish to accept the default value, type “yes” and press <Enter>.</p> <p>Otherwise, enter the desired box IPv4 netmask</p>
Step 10: HA Chassis IPv6 address? (Example: [fdfd:fdfd:7:145::1000:4814])	<p>Perform this step to configure the box IPv6.</p> <p>If you wish to accept the default value, type “yes” and press <Enter>.</p> <p>Otherwise, enter the desired box IPv6</p>
Step 11: HA Chassis Management IPv6 masklen? (Example: [33])	<p>Perform this step to configure the box IPv6 masklen.</p> <p>If you wish to accept the default value, type “yes” and press <Enter>.</p> <p>Otherwise, enter the desired box IPv6 masklen.</p>

Wizard Session Display (Example)	Comments
<pre> You have entered the following information: Hostname: <switch name> Use DHCP on mgmt0 interface: yes Enable IPv6: yes Enable IPv6 autoconfig (SLAAC) on mgmt0 interface: yes Enable DHCPv6 on mgmt0 interface: yes Admin password (Enter to leave unchanged): (CHANGED) HA Chassis IP address: 10.6.166.200 HA Chassis Management IP netmask: 255.255.255.0 HA Chassis IPv6 address: fdfd:fdfd:7:145::1000:4814 HA Chassis Management IPv6 masklen: 33 To change an answer, enter the step number to return to. Otherwise hit <enter> to save changes and exit. Choice: <Enter> Configuration changes saved. To return to the wizard from the CLI, enter the "configuration jump-start" command from configuration mode. Launching CLI... <switch name> [standalone: master] > </pre>	<p>The wizard displays a summary of your choices and then asks you to confirm the choices or to re-edit them.</p> <p>Either press <Enter> to save changes and exit, or enter the configuration step number that you wish to return to.</p> <p>To run the command "configuration jump-start" you must be in Config mode.</p>

Static IP configuration:

Wizard Session Display (Example)
<pre> Do you want to use the wizard for initial configuration? y Step 1: Hostname? [switch-112126] Step 2: Use DHCP on mgmt0 interface? [yes] n Step 3: Use zeroconf on mgmt0 interface? [no] Step 4: Primary IP address? 192.168.10.4 Mask length may not be zero if address is not zero (interface mgmt0) Step 5: Netmask? [0.0.0.0] 255.255.255.0 Step 6: Default gateway? 192.168.10.1 Step 7: Primary DNS server? Step 8: Domain name? Step 9: Enable IPv6? [yes] yes Step 10: Enable IPv6 autoconfig (SLAAC) on mgmt0 interface? [no] no Step 11: Update time? [yyyy/mm/dd hh:mm:ss] Step 12: Enable password hardening? [yes] yes Step 13: Admin password (Enter to leave unchanged)? You have entered the following information: Hostname: switch-112126 Use DHCP on mgmt0 interface: no Use zeroconf on mgmt0 interface: no Primary IP address: 192.168.10.4 Netmask: 255.255.255.0 Default gateway: 192.168.10.1 Primary DNS server: Domain name: Enable IPv6: yes Enable IPv6 autoconfig (SLAAC) on mgmt0 interface: no Update time: yyyy/mm/dd hh:mm:ss Enable password hardening: yes Admin password (Enter to leave unchanged): (unchanged) To change an answer, enter the step number to return to. Otherwise hit <enter> to save changes and exit. Choice: Configuration changes saved. To return to the wizard from the CLI, enter the "configuration jump-start" command from configure mode. Launching CLI... <hostname>[standalone: master] > </pre>

IP zeroconf configuration for modular switch systems:

Wizard Session Display (Example)

Configuration wizard

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

```
Step 1: Hostname? [switch-mgmt1]
Step 2: Use DHCP on mgmt0 interface? [yes]
Step 3: Enable IPv6? [yes]
Step 4: Enable IPv6 autoconfig (SLAAC) on mgmt0 interface? [no]
Step 5: Enable DHCPv6 on mgmt0 interface? [yes]
Step 6: Admin password (Enter to leave unchanged)?
Step 7: HA Chassis IP address: [10.6.166.200]
Step 8: HA Chassis Management IP netmask: [255.255.255.0]
Step 9: HA Chassis IPv6 address: [fdfd:fdfd:7:145::1000:4814]
Step 10: HA Chassis Management IPv6 masklen: [33]
```

You have entered the following information:

1. Hostname: sw-mantaray-201-mgmt1
2. Use DHCP on mgmt0 interface: yes
3. Enable IPv6: yes
4. Enable IPv6 autoconfig (SLAAC) on mgmt0 interface: no
5. Enable DHCPv6 on mgmt0 interface: yes
6. Admin password (Enter to leave unchanged): (unchanged)
7. HA Chassis IP address: 10.6.166.200
8. HA Chassis Management IP netmask: 255.255.255.0
9. HA Chassis IPv6 address: fdfd:fdfd:7:145::1000:4814
10. HA Chassis Management IPv6 masklen: 33

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

Choice:

Configuration changes saved.

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

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

```
switch # show interfaces mgmt0

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

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

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

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

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

2.1.1 Configuring the Switch with ZTP

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

For more information, please refer to section [“Zero-touch Provisioning”](#).

2.1.2 Rerunning the Wizard

To rerun the wizard:

1. Enter Config mode. Run:

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

2. Rerun the wizard. Run:

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

2.2 Starting the Command Line (CLI)

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

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

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

```
NVIDIA MLNX-OS Switch Management
Password:
Last login: <time> from <ip-address>

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

2.3 Starting the Web User Interface (WebUI)

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

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

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

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


In order to access WebUI through Safari 5.3, enable http:

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


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




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



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



- A default status summary is displayed.



2.4 Zero-touch Provisioning

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

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

Only a textual configuration file is supported.

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

Only one docker container can be launched in ZTP.

2.4.1 Running DHCP-ZTP

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

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

DHCPv4

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

DHCPv6

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

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

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

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

- System Model
- Chassis Part Number
- Chassis Serial Number
- Management MAC
- System Profile
- MLNX-OS Release Version

The corresponding subtypes respectively are defined as:

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

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

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

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

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

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

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

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

- a. Fetch the docker image file
- b. Load the docker image

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

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

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

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

2.4.2 ZTP on Modular Switches

For modular switch systems, the two management nodes start ZTP individually. Status synchronization is then performed between the two nodes:

- Target software image version needs to be the same, otherwise ZTP fails
- Both nodes must install the software image successfully, otherwise ZTP fails
- ZTP failure for one node leads to failure for both
- ZTP disable on one node leads to ZTP disable for both
- ZTP abort on one node leads to ZTP abort for both

In ZTP configuration files, commands between #<CHASSIS_MASTER> and #</CHASSIS_MASTER> pair are only executed on the master.

```
#<CHASSIS_MASTER>
chassis ha bip 10.7.146.34 /24
#</CHASSIS_MASTER>
```

Node reboot caused by ZTP is also synchronized:

1. Master node asks slave to reboot.
2. Slave node switches to next boot location and acknowledges the reboot request.
3. Master node reboots slave node via hardware.
4. Master node reboots itself.

2.4.3 ZTP and OS Upgrade

Software upgrade from non-ZTP versions to ZTP versions and vice versa is supported. When upgrading from a non-ZTP version, ZTP is disabled because ZTP is always assumed to start with an empty configuration, otherwise the final configuration becomes a mixture of the existing

configuration from the stored database and new configuration from the server and hence not deterministic.

2.4.4 DHCPv4 Configuration Example

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

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

DHCPv4 request is made out of the following components:

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

2.4.5 DHCPv6 Configuration Example

The following is a DHCPv6 configuration example:

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

DHCPv6 request is made out of the following components:

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

2.4.6 ZTP Commands

2.4.6.1 no zero-touch suppress-write

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

Notes	<ul style="list-style-type: none"> • When ZTP is active, “configuration write” is suppressed because it may interfere with ZTP operation. Therefore, after running “no zero-touch suppress-write” if “configuration write” is performed, then ZTP is disabled as a consequence of the database save. • To automatically save the configuration at the end of applying a configuration via ZTP, append the following two commands to the end of the config files. The first command will turn off the ZTP suppress-write, then the configuration write command should work. <ul style="list-style-type: none"> • no zero-touch suppress-write • configuration write
-------	---

2.4.6.2 zero-touch abort

	zero-touch abort Aborts on-going zero-touch process.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config
History	3.6.5000
Example	<pre>switch (config) # zero-touch abort Zero-touch failed [Zero-touch is aborted by operator] Zero-touch provisioning will be aborted</pre>
Related Commands	show zero-touch
Notes	

2.4.6.3 show zero-touch

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

2.5 Licenses

The 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.

2.5.1 Installing OS License via CLI

To install a license via CLI:

1. Before applying a license, please make sure your system's time is configured correctly by manually setting it using the CLI command "clock set", or by using NTP using the command "ntp".
2. Login as admin and change to Config mode.

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

3. Install the license using the key. Run:

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

4. Display the installed license(s) using the following command. Run:

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

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

5. Save the configuration to complete the license installation. Run:

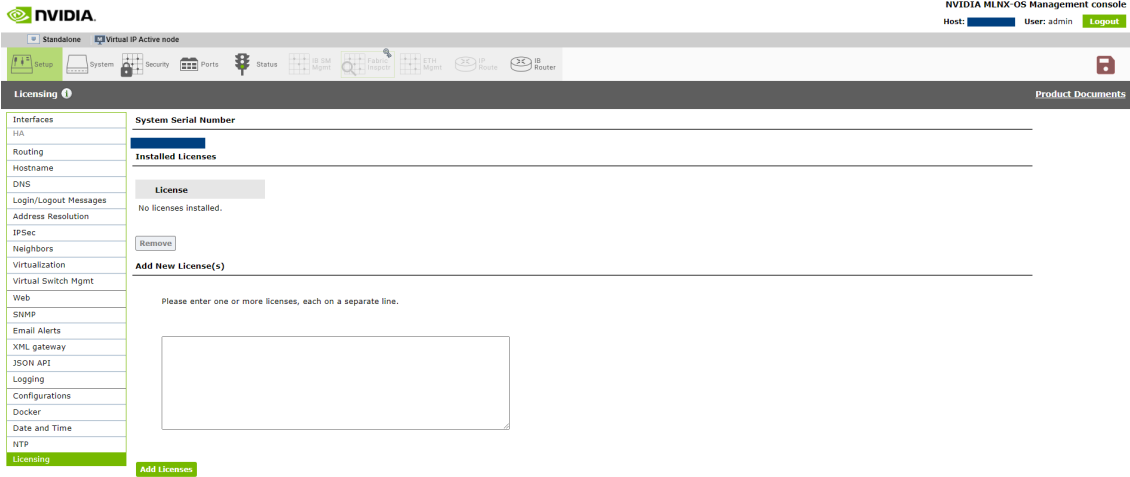
```
switch (config) # configuration write
```

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

2.5.2 Installing OS License via Web

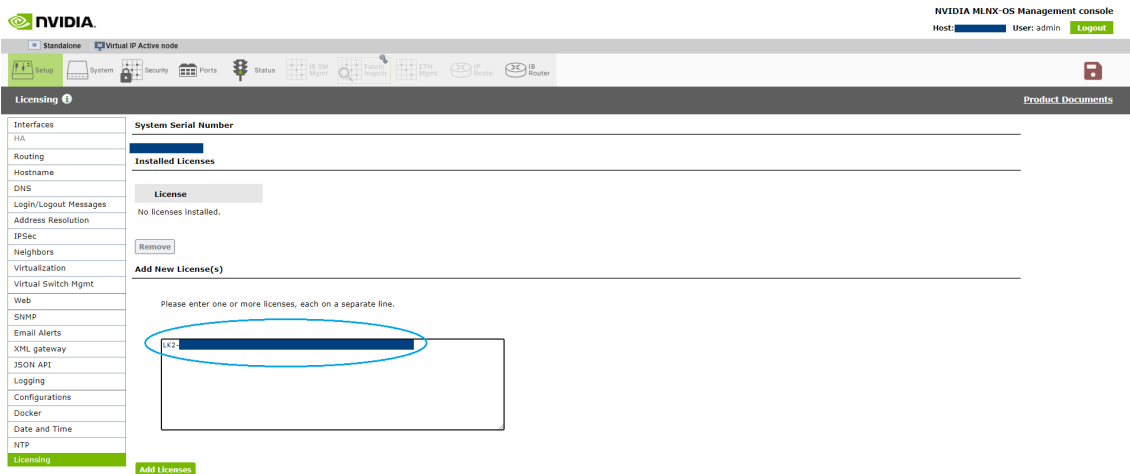
To install a license via WebUI:

1. Login as *admin*.
2. Click the Setup tab and then Licensing on the left side navigation pane.




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.



4. All installed licenses should now be displayed.



5. 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.

2.5.3 Retrieving a Lost License Key

In case of a lost license key, contact your authorized NVIDIA reseller and provide the switch's chassis serial number.

To obtain the switch's chassis serial number:

1. Log in to the switch.
2. Retrieve the switch's chassis serial number using the command "show inventory".

```
switch (config) # show inventory
-----
Module           Part Number      Serial Number     Asic Rev.   HW Rev.
-----
CHASSIS          MSB7800-ES2F     MT1602X17464     N/A         A1
MCMT             MSB7800-ES2F     MT1602X17464     0           A1
FAN1             MTEF-FANF-A      MT1602X16943     N/A         A3
FAN2             MTEF-FANF-A      MT1602X16944     N/A         A3
FAN3             MTEF-FANF-A      MT1602X16956     N/A         A3
FAN4             MTEF-FANF-A      MT1602X16957     N/A         A3
PS1              MTEF-PSF-AC-A    MT1601X09908     N/A         A3
```

3. Provide your authorized NVIDIA reseller with the chassis serial number for your system.
4. Once you receive the license key, you can install the license as described in the previous pages.

2.5.4 Additional Reading and Use Cases

For more information about getting started with NVIDIA Switches, please refer to the following Community post:

- [How To Get Started with NVIDIA Switches](#)

2.5.5 License Commands

- [2.1 Configuring the Switch for the First Time](#)
 - [2.1.1 Configuring the Switch with ZTP](#)
 - [2.1.2 Rerunning the Wizard](#)
- [2.2 Starting the Command Line \(CLI\)](#)
- [2.3 Starting the Web User Interface \(WebUI\)](#)
- [2.4 Zero-touch Provisioning](#)
 - [2.4.1 Running DHCP-ZTP](#)
 - [2.4.2 ZTP on Modular Switches](#)
 - [2.4.3 ZTP and OS Upgrade](#)
 - [2.4.4 DHCPv4 Configuration Example](#)
 - [2.4.5 DHCPv6 Configuration Example](#)
 - [2.4.6 ZTP Commands](#)
 - [2.4.6.1 no zero-touch suppress-write](#)
 - [2.4.6.2 zero-touch abort](#)
 - [2.4.6.3 show zero-touch](#)
- [2.5 Licenses](#)
 - [2.5.1 Installing OS License via CLI](#)
 - [2.5.2 Installing OS License via Web](#)
 - [2.5.3 Retrieving a Lost License Key](#)
 - [2.5.4 Additional Reading and Use Cases](#)
 - [2.5.5 License Commands](#)
 - [2.5.5.1 license delete](#)
 - [2.5.5.2 license install](#)
 - [2.5.5.3 show licenses](#)

2.5.5.1 license delete

	license delete <license-number> Removes license keys by ID.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.4.1100
Example	switch (config) # license delete <license-number>
Related Commands	license install show licenses
Notes	Before deleting a license from a switch which is configured to a system profile other than its default, the user must first disable all interfaces and then return the switch to its default system profile.

2.5.5.2 license install

	<code>license install<license-number></code> Installs a new license key.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.4.1100
Example	<code>switch (config) # licenses install <license-key></code>
Related Commands	<code>license delete</code> <code>show licenses</code>
Notes	

2.5.5.3 show licenses

	<code>show licenses</code> Displays a list of all installed licenses.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.4.1100
Example	<code>switch (config) # show licenses</code> License 1: <license key> Feature: SX_CONFIG Valid: yes Active: yes
Related Commands	<code>license delete</code> <code>license install</code>
Notes	For each license, the following is displayed: <ul style="list-style-type: none">• 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

3 User Interfaces

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

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

3.1 LED Indicators

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

- [SN2000 system LED indicators](#)
- [SN3000 system LED indicators](#)
- [SN4000 system LED indicators](#)

3.2 Command Line Interface (CLI)



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

3.2.1 CLI Modes

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

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

Configuration Mode	Description
Config interface management	Config Interface Management mode is a configuration mode for management interface mgmt0, mgmt1, and loopback.
Any command mode	Several commands, such as “show” commands, can be applied within any context.

3.2.2 Syntax Conventions

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

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

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

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

3.2.3 Getting Help

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

For example, typing “?” in the command line in Standard mode, will provide a link of the following available commands.

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

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

Try the following, to get started:

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

Enter “help” to view a description of the interactive help system.

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

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

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

3.2.4 Prompt and Response Conventions

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

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

The following session shows how to move between command modes:

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

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

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

3.2.5 Using the “no” Command Form

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

The command sequence below performs the following:

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

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

3.2.6 Parameter Key

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

Parameter	Description
<domain>	A domain name
<hostname>	A hostname (e.g., “switch-1”)
<ifname>	An interface name (e.g., “mgmt0”, “mgmt1”, “lo” (loopback), and so forth).
<index>	A number to be associated with aliased (secondary) IP addresses.
<IP address>	An IPv4 address (e.g., “192.168.0.1”)
<log level>	A syslog logging severity level. Possible values, from least to most severe, are as follows: “debug”, “info”, “notice”, “warning”, “error”, “crit”, “alert”, “emerg”.

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

3.2.7 CLI Pipeline Operator Commands

3.2.7.1 CLI Filtration Options “include” and “exclude”

The 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 (“|”). Quotation marks may be used to include or exclude a string including space, and multiple filters can be used simultaneously as shown in the example below.

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

Example:

```

switch (config) # show asic-version | include SIB2
MGMT          SIB2          15.2008.0236

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

```

3.2.7.2 CLI Monitoring Option “watch”

MLNX-OS

```

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

```

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

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

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

```

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

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

```

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

To exit “watch” mode, press Ctrl+C.

The “watch” option may be used in conjunction with the “include” and “exclude” options as follows:

```

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

```

Example:

```

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

```

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


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

Example:

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

3.2.7.3 CLI “json-print” Option

The MLNX-OS CLI supports printing “show” commands in JSON syntax.

To print the output of the “show” commands as JSON, use the following format:

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

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

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

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

For more information on JSON usage, please refer to [“JSON API”](#).

3.2.7.4 CLI Shortcuts

The following table presents the available keyboard shortcuts on the MLNX-OS CLI.

Key Combination	Description
Ctrl-a	Move cursor to beginning of line
Ctrl-b	Move cursor backward one character without deleting
Ctrl-c	Terminate operation
Ctrl-d	If cursor is in the middle of the line, delete one character forward If cursor is at the end of the line, show autocomplete options for current word or word fragment If cursor at an empty line, same as Esc
Ctrl-e	Move cursor to end of line
Ctrl-f	Move cursor forward one character
Ctrl-h	Delete one character backwards from cursor
Ctrl-i	Auto-complete current word (same as TAB)

Key Combination	Description
Ctrl-j	Return carriage (same as ENTER)
Ctrl-k	Delete line after cursor
Ctrl-l	Clear screen and show line at the top of terminal window
Ctrl-m	Return carriage (same as ENTER)
Ctrl-n	Next line (same as DOWN ARROW)
Ctrl-p	Next line (same as UP ARROW)
Ctrl-t	Transpose the two characters on either side of cursor
Ctrl-u	Delete line
Ctrl-w	Delete the last word
Ctrl-y	Retrieve (“yank”) last item deleted
Esc b	Move cursor one word backward
Esc c	Capitalizes first letter in word after cursor
Esc d	Delete one word forward from cursor
Esc f	Move one word forward from cursor
Esc l	Change word after cursor to lowercase letters
Esc Ctrl-h	Delete one word backward from cursor
Esc [A	Next line (same as DOWN ARROW)
Esc [B	Next line (same as UP ARROW)
Esc [C	Move forward one character from cursor
Esc [D	Move backward one character from cursor

3.3 Secure Shell (SSH)



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

3.3.1 Adding a Host and Providing an SSH Key

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

1. Change to Config mode.

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

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

```
switch (config) # ssh client global known-host "myserver ssh-rsa
AAAAB3NzaClyc2EAAAABIAAAIEAsXeklqc8T0EN2mnMcVcfhueaRYzIVqt4rVsrERIjmlJh4mkYYIa8hGGikNa+t5xw2dRrNxnHYLK51bU
SSG1ZNwZT1Dpme3pAZeMY7G4ZMgGIW9xOuaXgAA3eBeoUjFdi6+1BqchWk0nTb+gmfI/MK/heQNns7AtTrvgg/O5ryIc="
```

3. Verify what keys exist in the host.

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

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

No SSH user identities configured.
No SSH authorized keys configured.
```

RSA2 and a DSA2 host keys are generated by default. The RSA2 key can be used as SSH server and client, while DSA2 key can only be used as SSH client.

When the switch is a server, use RSA key to connect to the device.

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

3.3.2 Retrieving Return Codes When Executing Remote Commands

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

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

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

3.4 Web Interface Overview

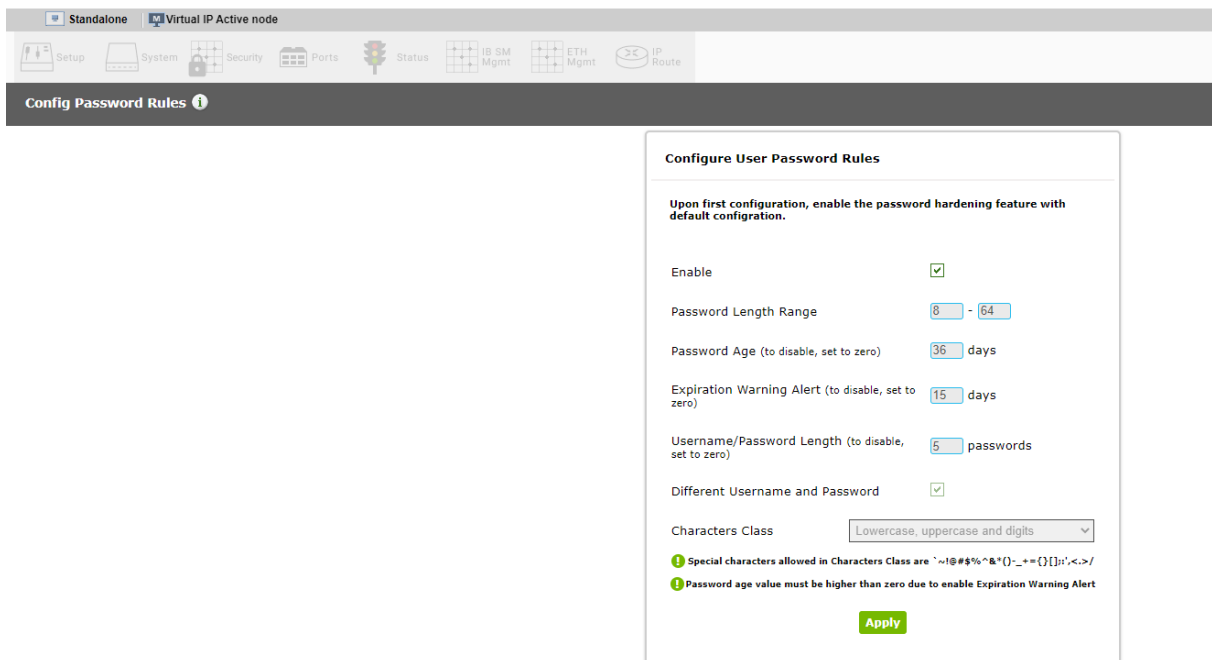


The MLNX-OS package equipped with web-based GUI that accepts input and provides output by generating webpages that can be viewed by the user using a web browser.

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

3.4.1 Password Hardening

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

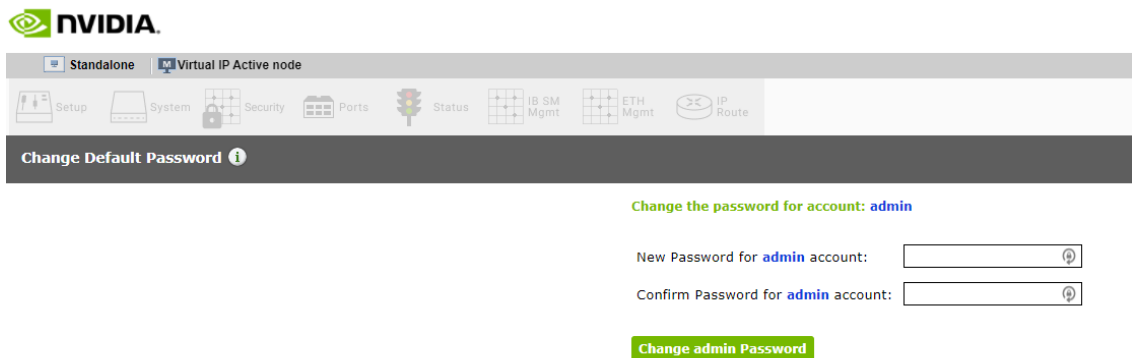


3.4.2 Changing Default Password

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


Upon initial login do the following:

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




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

locked.



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



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

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

3.4.3 About Web UI

The web interface makes available the following perspective tabs:

- Setup
- System
- Security
- Ports
- Status
- IB SM Management
- IB Router

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

Ports

Splitter ports: 1/1 1/2

Port Info	Port Counters
Port number : 1/1	RX bytes : 0
Port type : IB	RX packets : 0
IB Subnet : infiniband-default	RX errors : 0
Port description :	Symbol errors : 0
Logical port state : Down	VL15 dropped packets : 0
Physical port state : Polling	
Current line rate : -	
Supported speeds : sdr, qdr, fdr, edr, hdr, ndr	
Speed : -	
Supported widths : 1X, 2X, 4X	
Width : 4X	
Max supported MTUs : 4222	
MTU : 0	
VL admin capabilities : VL0 - VL7	
Operational VLS : -	
Threshold Level : N/A	
	TX bytes : 0
	TX packets : 0
	TX wait : 0
	TX discarded packets : 0

Clear Port 1/1 Counters

3.4.4 Setup Menu

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

Submenu Title	Description
Interfaces	Obtains the status of, configures, or disables interfaces to the fabric. Thus, you can: set or clear the IP address and netmask of an interface; enable DHCP to dynamically assign the IP address and netmask; and set interface attributes such as MTU, speed, duplex, etc.
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
Address Resolution	Adds static and dynamic ARP entries, and clears the dynamic ARP cache
IPSec	Configures IPSec
Neighbors	Displays IPv6 neighbor discovery protocol
Virtualization	Manages the virtualization and virtual machines
Virtual Switch Mgmt	Configures the system profile
Web	Configures web user interface and proxy settings
SNMP	Configures SNMP attributes, SNMP admin user, and trap sinks

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

3.4.5 System Menu

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

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

3.4.6 Security Menu

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

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

3.4.7 Ports Menu

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

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

3.4.8 Status Menu

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

Submenu Title	Description
Summary	Displays general information about the switch system and the OS image, including current date and time, hostname, uptime of system, system memory, CPU load averages, etc.
Profile and Capabilities	Displays general information about the switch system capabilities such as the enabled profiles (e.g IB/ETH) and their corresponding values
What Just Happened	Displays and configures What Just Happened packet drop reasons

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

3.4.9 IB SM Mgmt Menu

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

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

3.4.10 IB Router Menu

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

Submenu Title	Description
IB Router Global	Enables/disables IB router
IB Router Configuration	Manages IB router admin state and IB router interfaces

3.5 UI Commands



3.5.1 CLI Session

- [3.5.1 CLI Session](#)
 - [3.5.1.1 cli clear-history](#)
 - [3.5.1.2 cli default](#)
 - [3.5.1.3 cli max-sessions](#)
 - [3.5.1.4 cli session](#)
 - [3.5.1.5 terminal](#)
 - [3.5.1.6 terminal sysrq enable](#)
 - [3.5.1.7 show cli](#)
 - [3.5.1.8 show cli max-sessions](#)
 - [3.5.1.9 show cli num-sessions](#)
 - [3.5.1.10 Banner](#)
 - [3.5.1.10.1 banner login](#)
 - [3.5.1.10.2 banner login-local](#)
 - [3.5.1.10.3 banner login-remote](#)
 - [3.5.1.10.4 banner logout](#)
 - [3.5.1.10.5 banner logout-local](#)
 - [3.5.1.10.6 banner logout-remote](#)
 - [3.5.1.10.7 banner motd](#)
 - [3.5.1.10.8 show banner](#)
 - [3.5.1.11 SSH](#)
 - [3.5.1.11.1 ssh server enable](#)
 - [3.5.1.11.2 ssh server host-key](#)
 - [3.5.1.11.3 ssh server listen](#)
 - [3.5.1.11.4 ssh server login attempts](#)
 - [3.5.1.11.5 ssh server login timeout](#)
 - [3.5.1.11.6 ssh server login record-period](#)
 - [3.5.1.11.7 ssh server min-version](#)
 - [3.5.1.11.8 ssh server ports](#)
 - [3.5.1.11.9 ssh server security strict](#)
 - [3.5.1.11.9 ssh server security strict](#)

- [3.5.1.11.11 ssh server x11-forwarding](#)
- [3.5.1.11.12 ssh client global](#)
- [3.5.1.11.13 ssh client user](#)
- [3.5.1.11.14 slogin](#)
- [3.5.1.11.15 show ssh client](#)
- [3.5.1.11.16 show ssh server](#)
- [3.5.1.11.17 show ssh server host-keys](#)
- [3.5.1.11.18 show ssh server login record-period](#)
- [3.5.1.12 Remote Login](#)
 - [3.5.1.12.1 telnet](#)
 - [3.5.1.12.2 telnet-server enable](#)
 - [3.5.1.12.3 show telnet-server](#)
- [3.5.2 Web Interface](#)
 - [3.5.2.1 web auto-logout](#)
 - [3.5.2.2 web cache-enable](#)
 - [3.5.2.3 web client cert-verify](#)
 - [3.5.2.4 web client ca-list](#)
 - [3.5.2.5 web enable](#)
 - [3.5.2.6 web http](#)
 - [3.5.2.7 web httpd](#)
 - [3.5.2.8 web https](#)
 - [3.5.2.9 web https ssl renegotiation enable](#)
 - [3.5.2.10 web https ssl secure-cookie enable](#)
 - [3.5.2.11 web proxy auth authtype](#)
 - [3.5.2.12 web proxy auth basic](#)
 - [3.5.2.13 web session timeout](#)
 - [3.5.2.14 web session renewal](#)
 - [3.5.2.15 show web](#)

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

3.5.1.1 cli clear-history

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

3.5.1.2 cli default

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

3.5.1.3 cli max-sessions

	<p>cli max-sessions <number></p> <p>no cli max-sessions</p> <p>Configures the maximum number of simultaneous CLI sessions allowed. The no form of the command resets this value to its default.</p>	
Syntax Description	number	Range: 3-30
Default	30 sessions	
Configuration Mode	config	
History	3.5.0200	

Example	<code>switch (config) # cli max-sessions 40</code>
Related Commands	<code>show terminal</code>
Notes	

3.5.1.4 cli session

	<code>cli session {auto-logout <minutes> paging enable prefix-modes enable progress enable terminal {length <size> resize type <terminal-type> width} x-display full <display>}</code> <code>no cli session {auto-logout paging enable prefix-modes enable progress enable terminal type x-display}</code> Configures CLI options for this session only. The no form of the command deletes or disables the CLI sessions.	
Syntax Description	minutes	Configures keyboard inactivity timeout for automatic logout. Range: 0-35791 minutes Setting the value to 0 or using the no form of the command disables the auto logout.
	paging enable	Enables text viewing one screen at a time.
	prefix-modes enable	Configures the prefix modes feature of CLI and enables prefix modes for current session.
	progress enable	Enables progress updates.
	terminal length	Sets the number of lines for the current terminal. Range: 5-999
	terminal resize	Resizes the CLI terminal settings (to match the actual terminal window).
	terminal-type	Sets terminal type. Valid options are: <ul style="list-style-type: none"> • ansi • console • dumb • linux • unknown • vt52 • vt100 • vt102 • vt220 • xterm
	terminal width	Sets the width of the terminal in characters. Range: 34-999
	x-display full <display>	Specifies the display as a raw string (e.g. localhost:0.0).
Default	N/A	
Configuration Mode	config	
History	3.1.0000 3.8.2100: Removed "prefix-modes show-config" option and terminal type vt320	
Example	<code>switch (config) # cli session auto-logout</code>	
Related Commands	<code>show terminal</code>	
Notes	The "minutes" attribute can be configured from the CLI shell only.	

3.5.1.5 terminal

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

3.5.1.6 terminal sysrq enable

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

3.5.1.7 show cli

	show cli Displays the CLI configuration and status.	
Syntax Description	N/A	

Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<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 current session: Auto-logout: disabled Paging: enabled Progress tracking: enabled Prefix modes: enabled (and use in 'show configuration') Settings for current session: Show hidden config: yes Confirm losing changes: yes Confirm reboot/shutdown: no Confirm factory reset: yes Prompt on empty password: yes</pre>
Related Commands	cli default
Notes	

3.5.1.8 show cli max-sessions

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

3.5.1.9 show cli num-sessions

	<pre>show cli num-sessions Displays current number of sessions.</pre>
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.5.0200

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

3.5.1.10 Banner

3.5.1.10.1 banner login

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

3.5.1.10.2 banner login-local

	banner login-local <string> no banner login-local Sets system login local banner. The no form of the command resets the banner to its default value.
Syntax Description	N/A
Default	""
Configuration Mode	Any command mode
History	3.1.0000 3.5.0200: Added the no form of the command
Example	switch (config) # banner login-local Example
Related Commands	show banner
Notes	<ul style="list-style-type: none"> The login-local refers to the serial connection banner If more than one word is used (there is a space) quotation marks should be added (i.e., "xxxx xxxx")

3.5.1.10.3 banner login-remote

	banner login-remote <string> no banner login-remote Sets system login remote banner. The no form of the command resets the banner to its default value.	
Syntax Description	string	Text string
Default	""	
Configuration Mode	config	
History	3.1.0000 3.5.0200: Added the no form of the command	
Example	switch (config) # banner login-remote Example	
Related Commands	show banner	
Notes	<ul style="list-style-type: none"> The login-remote refers to the SSH connections banner If more than one word is used (there is a space) quotation marks should be added (i.e., "xxxx xxxx"). 	

3.5.1.10.4 banner logout

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

3.5.1.10.5 banner logout-local

	banner logout-local <string> no banner logout-local Sets system logout local banner. The no form of the command resets the banner to its default value.	
Syntax Description	string	Text string
Default	""	
Configuration Mode	config	
History	3.5.0200	

Example	<code>switch (config) # banner logout-local Example</code>
Related Commands	<code>show banner</code>
Notes	<ul style="list-style-type: none"> • The <code>logout-local</code> refers to the serial connection banner • If more than one word is used (there is a space) quotation marks should be added (i.e., <code>"xxxx xxxx"</code>).

3.5.1.10.6 banner logout-remote

	<code>banner logout-remote <string></code> <code>no banner logout-remote</code> Sets system logout remote banner. The no form of the command resets the banner to its default value.	
Syntax Description	string	Text string
Default	""	
Configuration Mode	config	
History	3.5.0200	
Example	<code>switch (config) # banner logout-remote Example</code>	
Related Commands	<code>show banner</code>	
Notes	<ul style="list-style-type: none"> • The <code>logout-remote</code> refers to SSH connections banner • If more than one word is used (there is a space) quotation marks should be added (i.e., <code>"xxxx xxxx"</code>). 	

3.5.1.10.7 banner motd

	<code>banner motd <string></code> <code>no banner motd</code> Configures the message of the day banner. The no form of the command resets the system Message of the Day banner.	
Syntax Description	string	Text string
Default	NVIDIA Switch	
Configuration Mode	config	
History	3.1.0000	
Example	<code>switch (config) # banner motd "My Banner"</code>	
Related Commands	<code>show banner</code>	
Notes	<ul style="list-style-type: none"> • If more than one word is used (there is a space) quotation marks should be added (i.e., <code>"xxxx xxxx"</code>). • To insert a multi-line MotD, hit Ctrl-V (escape sequence) followed by Ctrl-J (new line sequence). The symbol <code>“^J”</code> should appear. Then, whatever is typed after it becomes the new line of the MotD. Remember to also include the string between quotation marks. 	

3.5.1.10.8 show banner

	show banner Sets system logout remote banner. The no form of the command resets the banner to its default value.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.5.0200	Updated example
	3.6.6000	Updated example
	3.9.3200	Updated example
Example	<pre>switch (config) # show banner Banners: Message of the Day (MOTD): Switch Login: NVIDIA MLNX-OS Switch Management Logout: Goodbye</pre>	
Related Commands	banner login banner login-local banner login-remote banner logout banner logout-local banner logout-remote banner motd	
Notes		

3.5.1.11 SSH

3.5.1.11.1 ssh server enable

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

3.5.1.11.2 ssh server host-key

	ssh server host-key {<key-type> {private-key <private-key> public-key <public-key>} generate} Configures host keys for SSH.	
Syntax Description	key-type	<ul style="list-style-type: none"> rsa2—RSAv2 dsa2—DSA v2
	private-key	Sets new private-key for the host keys of the specified type.
	public-key	Sets new public-key for the host keys of the specified type.
	generate	Generates new RSA and DSA host keys for SSH.
Default	SSH keys are locally generated	
Configuration Mode	config	
History	3.1.0000 3.4.2300: Added notes 3.9.0300: Removed RSAv1 3.9.1000: Added a note	
Example	<pre>switch (config) # ssh server host-key dsa2 private-key Key: ***** Confirm: *****</pre>	
Related Commands	show banner	
Notes	<p>RSA2 and a DSA2 host keys are generated by default. The RSA2 key can be used as SSH server and client, while DSA2 key can only be used as SSH client.</p> <p>When the switch is a server, use RSA key to connect to the NVIDIA Onyx device.</p> <p>When the switch is a client (e.g. downloading image or uploading logs), RSA key is recommended. DSA key is only for legacy devices and has been deprecated by OpenSSH starting with the 7.0 release.</p>	

3.5.1.11.3 ssh server listen

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

3.5.1.11.4 ssh server login attempts

	ssh server login attempts <number> no ssh server login attempts Configures maximum login attempts on SSH server. The no form of the command resets the login attempts value to its default.	
Syntax Description	number	Range: 3-100 attempts
	interface	Adds interface to SSH server access restriction list. Possible interfaces are "lo", and "mgmt0".
Default	6 attempts	
Configuration Mode	config	
History	3.1.0000 3.5.1000: Increased minimum number of attempts 3.9.0900: Added notes	
Example	switch (config) # ssh server login attempts 5	
Related Commands	show ssh server	
Notes	<ul style="list-style-type: none"> • The number configured with this command will be relevant only if it is equal or smaller than the number of password prompts • Be aware that the "aaa authentication attempts lockout max-fail" default is 5, and the user might be locked before this command will have an affect. Both numbers need to be configured 	

3.5.1.11.5 ssh server login timeout

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

3.5.1.11.6 ssh server login record-period

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

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

3.5.1.11.7 ssh server min-version

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

3.5.1.11.8 ssh server ports

	ssh server ports {<port1> [<port2>...]} Specifies which ports the SSH server listens on.	
Syntax Description	port	Port number between [1-65535]
Default	22	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ssh server ports 22	
Related Commands	show ssh server	
Notes	<ul style="list-style-type: none"> Multiple ports can be specified by repeating the <port> parameter The command will remove any previous ports if not listed in the command 	

3.5.1.11.9 ssh server security strict

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

Configuration Mode	config
History	3.3.5060 3.6.4000 3.9.0300: Updated notes
Example	switch (config) # ssh server security strict
Related Commands	show ssh server
Notes	The following ciphers are disabled for SSH when strict security is enabled: <ul style="list-style-type: none"> • aes256-cbc • aes192-cbc • aes128-cbc • rijndael-cbc@lysator.liu.se • 3des-cbc

3.5.1.11.10 ssh server security strict

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

3.5.1.11.11 ssh server x11-forwarding

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

3.5.1.11.12 ssh client global

	<code>ssh client global {host-key-check <policy>} known-host <known-host-entry>}</code> <code>no ssh client global {host-key-check known-host localhost}</code> Configures global SSH client settings. The no form of the command negates global SSH client settings.	
Syntax Description	<code>host-key-check <policy></code>	Sets SSH client configuration to control how host key checking is performed. This parameter may be set in 3 ways. <ul style="list-style-type: none"> • 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
	<code>known-host</code>	Adds an entry to the global known-hosts configuration file
	<code>known-host-entry</code>	Adds/removes an entry to/from the global known-hosts configuration file. The entry consist of “<IP> <key-type> <key>”.
Default	host-key-check - ask, no keys are configured by default	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # ssh client global host-key-check no switch (config) # ssh client global known-host "72.30.2.2 ssh-rsa AAAAB3NzaC1yc2EAAAAB...f2CyXFq4pzaR1jar1Vk="</pre>	
Related Commands	show ssh client	
Notes		

3.5.1.11.13 ssh client user

	<code>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}</code> <code>no ssh client user admin {authorized-key sshv2 <public key ID> identity <key type>}</code> 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	<code>username</code>	The specified user must be a valid account on the system. Possible values for this parameter are “admin”, “monitor”, “xmladmin”, and “xmluser”.
	<code>authorized-key sshv2 <public key></code>	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.
	<code>identity <key type></code>	Sets certain SSH client identity settings for a user, dsa2 or rsa2.
	<code>generate</code>	Generates SSH client identity keys for specified user.
	<code>private-key</code>	Sets private key SSH client identity settings for the user.
	<code>public-key</code>	Sets public key SSH client identity settings for the user.

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

3.5.1.11.14 slogin

	slogin [<slogin options>] <hostname> Invokes the SSH client. The user is returned to the CLI when SSH finishes.		
Syntax Description	slogin options	-p -c -L -l -m -R -o -1 -2 -4 -6 -g -q -V -v -x -X -Y -y -a -A	-o flags (option allowed flags): AdressFamily BatchMode CheckHostIP Cipher Ciphers ConnectTimeout ForwardAgent ForwardX11 ForwardX11Trusted HostKeyAlgorithms KexAlgorithms LogLevel MACs Port PubkeyAcceptedKeyTypes PubkeyAuthentication StrictHostKeyChecking TCPKeepAlive User VerifyHostKeyDNS
	vrf_name	There are no restrictions on the VRF name, as long as the VRF exists in the switch.	
Default	N/A		
Configuration Mode	config		
History	3.1.0000 3.10.1000: Updated the slogin options		
Example	switch (config) # slogin 192.168.10.70 The authenticity of host '192.168.10.70 (192.168.10.70)' can't be established. RSA key fingerprint is 2e:ad:2d:23:45:4e:47:e0:2c:ae:8c:34:f0:1a:88:cb. Are you sure you want to continue connecting (yes/no)? yes		
Related Commands			

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

3.5.1.11.15 show ssh client

	<code>show ssh client</code> Displays the client configuration of the SSH server.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<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.</pre>
Related Commands	
Notes	

3.5.1.11.16 show ssh server

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

Example	<pre>switch (config) # show ssh server SSH server configuration: SSH server enabled: yes Server security strict mode: no Minimum protocol version: 2 TCP forwarding enabled: yes X11 forwarding enabled: no SSH login timeout: 120 SSH login max attempts: 6 SSH server login record-period: 1 SSH server ports: 22 Interface listen enabled: yes Listen Interfaces: No interface configured. Host Key Finger Prints and Key Lengths: RSA v2 host key: SHA256:gVu6qLW1ZifEp8wRer2jvkILZMGN16VCYU3HqC1INC8 (2048) DSA v2 host key: SHA256:JnldTEla20ZF/c5LdIqo9251DzO742k3hFCQh3Jt4ZA (1024)</pre>
Related Commands	
Notes	

3.5.1.11.17 show ssh server host-keys

	<pre>show ssh server host-keys Displays SSH host key configuration.</pre>
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	<pre>3.1.0000 3.6.6000: Updated example 3.9.0300: Updated example—removed RSA v1</pre>
Example	<pre>switch (config) # show ssh server host-keys SSH server configuration: SSH server enabled: yes Server security strict mode: no Minimum protocol version: 2 TCP forwarding enabled: yes X11 forwarding enabled: no SSH login timeout: 120 SSH login max attempts: 6 SSH server ports: 22 Interface listen enabled: yes Listen Interfaces: No interface configured. Host Key Finger Prints and Key Lengths: RSA v2 host key: SHA256:gVu6qLWLZifEp8wRer2jvkILZMGN16VCYU3HqC1INC8 (2048) DSA v2 host key: SHA256:JnldTEla20ZF/c5LdIqo9251DzO742k3hFCQh3Jt4ZA (1024) Host Keys: RSA v2 host key: "kebo-2100-1 ssh-rsa AAAAB3Nza<...>KE5" DSA v2 host key: "kebo-2100-1 ssh-dss AAAAB3Nza<...>/s="</pre>
Related Commands	ssh server host-keys
Notes	

3.5.1.11.18 show ssh server login record-period

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

3.5.1.12 Remote Login

3.5.1.12.1 telnet

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

3.5.1.12.2 telnet-server enable

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

Related Commands	telnet-server show telnet-server
Notes	

3.5.1.12.3 show telnet-server

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

3.5.2 Web Interface

3.5.2.1 web auto-logout

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

3.5.2.2 web cache-enable

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

3.5.2.3 web client cert-verify

	<p>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.</p>
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.2.3000
Example	switch (config) # web client cert-verify
Related Commands	
Notes	

3.5.2.4 web client ca-list

	<p>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.</p>	
Syntax Description	ca-list-name	Specifies CA list to configure
	default-ca-list	Configures default supplemental CA certificate list
	none	Uses no supplemental certificates
Default	default-ca-list	
Configuration Mode	config	
History	3.2.3000	

Example	<code>switch (config) # web client ca-list default-ca-list</code>
Related Commands	
Notes	

3.5.2.5 web enable

	<code>web enable</code> <code>no web enable</code> 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 3.8.1000—Added note
Example	<code>switch (config) # web enable</code>
Related Commands	<code>show web</code>
Notes	Disabling WebUI or HTTPS blocks connected LCD tablet display of CS8500 modular switch.

3.5.2.6 web http

	<code>web http {enable port <port-number> redirect}</code> <code>no web http {enable port redirect}</code> 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	<ul style="list-style-type: none"> • HTTP is disabled • HTTP TCP port is 80 • HTTP redirect to HTTPS is disabled 	
Configuration Mode	config	
History	3.1.0000	
Example	<code>switch (config) # web http enable</code>	
Related Commands	<code>show web</code> <code>web enable</code>	
Notes	Enabling HTTP is meaningful if the WebUI as a whole is enabled	

3.5.2.7 web httpd

	web httpd listen {enable interface <ifName>} no web httpd listen {enable interface <ifName>} Enables the listen interface restricted list for HTTP and HTTPS. The no form of the command disables the HTTP server listen ability.	
Syntax Description	enable	Enables Web interface restrictions on access to this system.
	interface <ifName>	Adds interface to Web server access restriction list (i.e., mgmt0, mgmt1).
Default	<ul style="list-style-type: none"> Listening is enabled All interfaces are permitted. 	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # web httpd listen enable	
Related Commands	show web web enable	
Notes	If enabled, and if at least one of the interfaces listed is eligible to be a listen interface, then HTTP/HTTPS requests will only be accepted on those interfaces. Otherwise, HTTP/HTTPS requests are accepted on any interface.	

3.5.2.8 web https

	web https {certificate {regenerate name default-cert} enable port <port number> ssl ciphers {all TLS TLS1.2}} no web https {enable port <port number>} Configures HTTPS access to the web-based management console. The no form of the command negates HTTPS settings for the web-based management console.	
Syntax Description	certificate regenerate	Re-generates certificate to use for HTTPS connections
	certificate name	Configure the named certificate to be used for HTTPS connections
	certificate default-cert	Configure HTTPS to use the configured default certificate
	enable	Enables HTTPS access to the web-based management console
	port	Sets a TCP port for HTTPS access
	ssl ciphers {all TLS TLS1.2}	Sets ciphers to be used for HTTPS
Default	<ul style="list-style-type: none"> HTTPS is enabled Default port is 443 	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Added “ssl ciphers” parameter
	3.4.0010	Added TLS parameter to “ssl ciphers”
	3.8.1000	Added note

Example	<code>switch (config) # web https enable</code>
Related Commands	<code>show web</code> <code>web enable</code>
Notes	<ul style="list-style-type: none"> • Enabling HTTPS is meaningful if the WebUI as a whole is enabled • Disabling WebUI or HTTPS blocks connected LCD tablet display of CS8500 modular switch • See the command “crypto certificate default-cert name” for how to change the default certificate if inheriting the configured default certificate is preferred

3.5.2.9 web https ssl renegotiation enable

	<code>web https ssl renegotiation enable</code> <code>no web https ssl renegotiation enable</code> Enables SSL renegotiation flag in httpd web server. The no form of the command disables SSL renegotiation flag in httpd web server.
Syntax Description	N/A
Default	<ul style="list-style-type: none"> • HTTPS is enabled • Default port is 443
Configuration Mode	config
History	3.6.8008
Example	<code>switch (config) # web https ssl renegotiation enable</code>
Related Commands	<code>show web</code> <code>web enable</code>
Notes	

3.5.2.10 web https ssl secure-cookie enable

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

3.5.2.11 web proxy auth authtype

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

3.5.2.12 web proxy auth basic

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

3.5.2.13 web session timeout

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

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

3.5.2.14 web session renewal

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

3.5.2.15 show web

	show web Displays WebUI configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.6000 3.6.8008—Updated example
Example	<pre>switch (config) # show web Web User Interface: Web interface enabled: yes Web caching enabled: no HTTP enabled: no HTTP port: 80 HTTP redirect to HTTPS: no HTTPS enabled: yes HTTPS port: 443 HTTPS ssl-ciphers: TLS1.2 HTTPS ssl-renegotiation: no HTTPS ssl-secure-cookie: yes HTTPS certificate name: default-cert Listen enabled: yes Listen Interfaces: No interface configured. Inactivity timeout: 1 hr Session timeout: 2 hr 30 min Session renewal: 30 min Web file transfer proxy: Proxy enabled: no Web file transfer certificate authority: HTTPS server cert verify: yes HTTPS supplemental CA list: default-ca-list</pre>

Related Commands	web auto-logout web cache-enable web enable web http web httpd web https web https ssl renegotiation enable web https ssl secure-cookie enable web proxy auth authtype web proxy auth basic
Notes	

4 System Management

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

- [Management Interfaces](#)
- [Chassis Management](#)
- [UNBREAKABLE-LINK® Adapter and Switch Technology](#)
- [Upgrade/Downgrade Process](#)
- [Configuration Management](#)
- [mDNS](#)

4.1 Management Interfaces



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

4.1.1 Configuring Management Interfaces with Static IP Addresses

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

1. Enter Config configuration mode. Run:

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

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

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

3. Define your interfaces statically using the following command:

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

4.1.2 Configuring IPv6 Address on the Management Interface

1. Enable IPv6 on this interface.

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

2. Set the IPv6 address to be configured automatically.

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

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

In this case the serial connection should be used.

4.1.4 Default Gateway

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

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

4.1.5 Configuring Hostname via DHCP (DHCP Client Option 12)

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

To enable fetching hostname from DHCP server, run the following:

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

To disable fetching hostname from DHCP server, run the following:

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

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

4.1.6 Management Interface Commands



- [4.1.6.1 Interface](#)
 - [4.1.6.1.1 interface](#)
 - [4.1.6.1.2 ip address](#)
 - [4.1.6.1.3 ip default-gateway](#)
 - [4.1.6.1.4 alias](#)
 - [4.1.6.1.5 mtu](#)
 - [4.1.6.1.6 duplex](#)
 - [4.1.6.1.7 speed](#)
 - [4.1.6.1.8 dhcp](#)
 - [4.1.6.1.9 dhcp hostname](#)
 - [4.1.6.1.10 shutdown](#)
 - [4.1.6.1.11 zeroconf](#)
 - [4.1.6.1.12 comment](#)
 - [4.1.6.1.13 ipv6 enable](#)
 - [4.1.6.1.14 ipv6 address](#)
 - [4.1.6.1.15 ipv6 dhcp primary-intf](#)
 - [4.1.6.1.16 ipv6 dhcp stateless](#)
 - [4.1.6.1.17 ipv6 dhcp client enable](#)
 - [4.1.6.1.18 ipv6 dhcp client renew](#)
 - [4.1.6.1.19 show interfaces mgmt0](#)
 - [4.1.6.1.20 show interfaces mgmt0 brief](#)
 - [4.1.6.1.21 show interfaces mgmt0 configured](#)
- [4.1.6.2 Hostname Resolution](#)
 - [4.1.6.2.1 hostname](#)
 - [4.1.6.2.2 ip name-server](#)
 - [4.1.6.2.3 ip domain-list](#)
 - [4.1.6.2.4 ip/ipv6 host](#)
 - [4.1.6.2.5 ip/ipv6 map-hostname](#)
 - [4.1.6.2.6 show hosts](#)
- [4.1.6.3 Routing](#)
 - [4.1.6.3.1 IP route](#)
 - [4.1.6.3.2 ipv6 default-gateway](#)

- [4.1.6.3.3 show ip/ipv6 route](#)
- [4.1.6.3.4 show ipv6 default-gateway](#)
- [4.1.6.4 Network to Media Resolution \(ARP & NDP\)](#)
 - [4.1.6.4.1 ipv6 neighbor](#)
 - [4.1.6.4.2 clear ipv6 neighbors](#)
 - [4.1.6.4.3 show ipv6 neighbors](#)
- [4.1.6.5 DHCP](#)
 - [4.1.6.5.1 ip dhcp](#)
 - [4.1.6.5.2 show ip dhcp](#)
- [4.1.6.6 IP Diagnostic Tools](#)
 - [4.1.6.6.1 ping](#)
 - [4.1.6.6.2 traceroute](#)
 - [4.1.6.6.3 tcpdump](#)

4.1.6.1 Interface

4.1.6.1.1 interface

	interface {mgmt0 mgmt1 lo vlan<id> ib0} Enters a management interface context.}	
Syntax Description	mgmt0	Management port 0 (out of band).
	mgmt1	Management port 1 (out of band).
	lo	Loopback interface.
	vlan<id>	In-band management interface (e.g., vlan10).
	ib0	IPoIB in-band management.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config)# interface mgmt0 switch (config interface mgmt0)#	
Related Commands	show interfaces <ifname>	
Notes		

4.1.6.1.2 ip address

	ip address <IP address> <netmask> no ip address Sets the IP address and netmask of this interface. The no form of the command clears the IP address and netmask of this interface.	
Syntax Description	IP address	IPv4 address
	netmask	Subnet mask of IP address
Default	0.0.0.0/0	

Configuration Mode	config interface management
History	3.1.0000
Example	<code>switch (config interface mgmt0)# ip address 10.10.10.10 255.255.255.0</code>
Related Commands	show interfaces <ifname>
Notes	If DHCP is enabled on the specified interface, then the DHCP IP assignment will hold until DHCP is disabled

4.1.6.1.3 ip default-gateway

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

4.1.6.1.4 alias

	<code>alias <index> ip address < IP address> <netmask></code> <code>no alias <index></code> Adds an additional IP address to the specified interface. The secondary address will appear in the output of “show interface” under the data of the primary interface along with the alias. The no form of the command removes the secondary address to the specified interface.	
Syntax Description	index	A number that is to be aliased to (associated with) the secondary IP.
	IP address	Additional IP address.
	netmask	Subnet mask of the IP address.
Default	N/A	
Configuration Mode	config interface management	
History	3.1.0000	
Example	<code>switch (config interface mgmt0)# alias 2 ip address 9.9.9.9 255.255.255.255</code>	
Related Commands	show interfaces <ifname>	

Notes	<ul style="list-style-type: none"> • If DHCP is enabled on the specified interface, then the DHCP IP assignment will hold until DHCP is disabled • More than one additional IP address can be added to the interface
-------	--

4.1.6.1.5 mtu

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

4.1.6.1.6 duplex

	<code>duplex <duplex></code> <code>no duplex</code> Sets the interface duplex. The no form of the command resets the duplex setting for this interface to its default value.	
Syntax Description	<code>duplex</code>	Sets the duplex mode of the interface. The following are the possible values: <ul style="list-style-type: none"> • half–half duplex • full–full duplex • auto–auto duplex sensing (half or full)
Default	auto	
Configuration Mode	config interface management	
History	3.1.0000	
Example	<code>switch (config interface mgmt0)# duplex auto</code>	
Related Commands	<code>show interfaces <ifname></code>	
Notes	<ul style="list-style-type: none"> • Setting the duplex to “auto” also sets the speed to “auto” • Setting the duplex to one of the settings “half” or “full” also sets the speed to a manual setting which is determined by querying the interface to find out its current auto-detected state 	

4.1.6.1.7 speed

	speed <speed> no speed Sets the interface speed. The no form of the command resets the speed setting for this interface to its default value.	
Syntax Description	speed	Sets the speed of the interface. The following are the possible values: <ul style="list-style-type: none"> • 10—fixed to 10Mbps • 100—fixed to 1000Mbps • 1000—fixed to 1000Mbps • auto—auto speed sensing (10/100/1000Mbps)
Default	auto	
Configuration Mode	config interface management	
History	3.1.0000	
Example	<pre>switch (config interface mgmt0)# speed auto</pre>	
Related Commands	show interfaces <ifname>	
Notes	<ul style="list-style-type: none"> • Setting the speed to “auto” also sets the duplex to “auto” • Setting the speed to one of the manual settings (generally “10”, “100”, or “1000”) also sets the duplex to a manual setting which is determined by querying the interface to find out its current auto-detected state 	

4.1.6.1.8 dhcp

	dhcp [renew] no dhcp Enables DHCP on the specified interface. The no form of the command disables DHCP on the specified interface.	
Syntax Description	renew	Forces a renewal of the IP address. A restart on the DHCP client for the specified interface will be issued.
Default	Could be enabled or disabled (per part number) manufactured with 3.2.0500	
Configuration Mode	config interface management	
History	3.1.0000 3.9.1900: Added note	
Example	<pre>switch (config interface mgmt0)# dhcp</pre>	
Related Commands	show interfaces <ifname> configured	
Notes	<ul style="list-style-type: none"> • When enabling DHCP, the IP address and netmask are received via DHCP hence, the static IP address configuration is ignored • Enabling DHCP disables zeroconf and vice versa • Setting a static IP address and netmask does not disable DHCP. DHCP is disabled using the “no” form of this command, or by enabling zeroconf. • When static IP is configured, DHCP will not run. 	

4.1.6.1.9 dhcp hostname

	<p>dhcp hostname no dhcp hostname Enables fetching the hostname from DHCP for this interface. The no form of the command disables fetching the hostname from DHCP for this interface.</p>
Syntax Description	N/A
Default	Enabled
Configuration Mode	config interface management
History	3.5.1000
Example	<code>switch (config interface mgmt0)# dhcp hostname</code>
Related Commands	<p>hostname <hostname> show interfaces <ifname> configured</p>
Notes	<ul style="list-style-type: none"> • If a hostname is configured manually by the user, that configuration would override the “dhcp hostname” configuration • When a default hostname is not configured, the DHCP server assigns the new hostname for your machine (after upgrading to version 3.5.1000) • These commands do not work on in-band interfaces

4.1.6.1.10 shutdown

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

4.1.6.1.11 zeroconf

	<p>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.</p>
Syntax Description	N/A
Default	no zeroconf

Configuration Mode	config interface management
History	3.1.0000
Example	<code>switch (config interface mgmt0)# zeroconf</code>
Related Commands	show interfaces <ifname> configured
Notes	Enabling zeroconf disables DHCP and vice versa.

4.1.6.1.12 comment

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

4.1.6.1.13 ipv6 enable

	ipv6 enable no ipv6 enable Enables all IPv6 addressing for this interface. The no form of the command disables all IPv6 addressing for this interface.	
Syntax Description	N/A	
Default	IPv6 addressing is disabled	
Configuration Mode	config interface management	
History	3.1.0000	
Example	<code>switch (config interface mgmt0)# ipv6 enable</code>	
Related Commands	ipv6 address show interface <ifname>	
Notes	<ul style="list-style-type: none"> The interface identifier is a 64-bit long modified EUI-64, which is based on the MAC address of the interface If IPv6 is enabled on an interface, the system will automatically add a link-local address to the interface. Link-local addresses can only be used to communicate with other hosts on the same link, and packets with link-local addresses are never forwarded by a router. A link-local address, which may not be removed, is required for proper IPv6 operation. The link-local addresses start with “fe80::”, and are combined with the interface identifier to form the complete address. 	

4.1.6.1.14 ipv6 address

	<pre>ipv6 address {<IPv6 address/netmask> autoconfig [default privacy]}</pre> <pre>no ipv6 {<IPv6 address/netmask> autoconfig [default privacy]}</pre> <p>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.</p>	
Syntax Description	IPv6 address/ netmask	Configures a static IPv6 address and netmask. Format example: 2001:db8:1234::5678/64.
	autoconfig	Enables IPv6 stateless address auto configuration (SLAAC) for this interface. An address will be automatically added to the interface based on an IPv6 prefix learned from router advertisements, combined with an interface identifier.
	autoconfig default	Enables default learning routes. The default route will be discovered automatically, if the autoconfig is enabled.
	autoconfig privacy	Uses privacy extensions for SLAAC to construct the autoconfig address, if the autoconfig is enabled.
Default	No IP address available, auto config is enabled	
Configuration Mode	config interface management	
History	3.1.0000	
Example	switch (config interface mgmt0)# ipv6 fe80::202:c9ff:fe5e:a5d8/64	
Related Commands	<pre>ipv6 enable</pre> <pre>show interface <ifname></pre>	
Notes	<ul style="list-style-type: none"> On a given interface, up to 16 addresses can be configured For Ethernet, the default interface identifier is a 64-bit long modified EUI-64, which is based on the MAC address of the interface 	

4.1.6.1.15 ipv6 dhcp primary-intf

	<pre>ipv6 dhcp primary-intf <if-name></pre> <pre>no ipv6 dhcp primary-intf</pre> <p>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.</p>	
Syntax Description	if-name	Interface name: <ul style="list-style-type: none"> lo mgmt0 mgmt1
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config)# ipv6 dhcp primary-intf mgmt0	
Related Commands	<pre>ipv6 enable</pre> <pre>ipv6 address</pre> <pre>show interface <ifname></pre>	
Notes		

4.1.6.1.16 ipv6 dhcp stateless

	<pre>ipv6 dhcp stateless no ipv6 dhcp stateless</pre> <p>Enables stateless DHCPv6 requests. The no form of the command disables stateless DHCPv6 requests.</p>
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Example	<pre>switch (config)# ipv6 dhcp stateless</pre>
Related Commands	<pre>ipv6 enable ipv6 address show interface <ifname></pre>
Notes	<ul style="list-style-type: none"> • This command only gets DNS configuration, not an IPv6 address • The no form of the command requests all information, including an IPv6 address

4.1.6.1.17 ipv6 dhcp client enable

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

4.1.6.1.18 ipv6 dhcp client renew

	<pre>ipv6 dhcp client renew</pre> <p>Renews DHCPv6 lease for this interface.</p>
Syntax Description	N/A
Default	N/A
Configuration Mode	config interface management
History	3.7.11xx
Example	<pre>switch (config interface mgmt0)# ipv6 dhcp client renew</pre>

Related Commands	ipv6 dhcp client enable show ipv6 dhcp
Notes	

4.1.6.1.19 show interfaces mgmt0

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

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

4.1.6.1.20 show interfaces mgmt0 brief

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

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

4.1.6.1.21 show interfaces mgmt0 configured

	<pre>show interface mgmt0 configured</pre> <p>Displays configuration information about the specified interface.</p>
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	<p>3.1.0000</p> <p>3.5.1000: Updated example with “DHCP Hostname”</p> <p>3.6.8008: Updated example</p>
Example	<pre>switch (config)# show interfaces mgmt0 configured Interface mgmt0 configuration: Comment : Enabled : yes DHCP : yes DHCP Hostname : yes Zeroconf : no IP address : Netmask : IPv6 enabled : yes Autoconf enabled: no Autoconf route : yes Autoconf privacy: no DHCPv6 enabled : yes IPv6 addresses : 0 Speed : auto Duplex : auto MTU : 1500</pre>
Related Commands	
Notes	

4.1.6.2 Hostname Resolution

4.1.6.2.1 hostname

	hostname <hostname> no hostname Sets a static system hostname. The no form of the command clears the system hostname.	
Syntax Description	hostname	A free-form string
Default	Default hostname	
Configuration Mode	config	
History	3.1.0000 3.6.3004: Added support for the character "."	
Example	switch (config)# hostname my-switch-hostname	
Related Commands	show hosts	
Notes	<ul style="list-style-type: none"> • Hostname may contain letters, numbers, periods ('.'), and hyphens ('-'), in any combination • Hostname may be 1-63 characters long • Hostname may not begin with a hyphen • Hostname may not contain other characters, such as "%", "_" etc. • Hostname may not be set to one of the valid logging commands (i.e. debug-files, fields, files, format, level, local, monitor, receive, trap) • Changing the hostname stamps a new HTTPS certificate 	

4.1.6.2.2 ip name-server

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

4.1.6.2.3 ip domain-list

	ip domain-list <domain-name> no ip domain-list <domain-name> Sets the static domain name. The no form of the command clears the domain name.	
--	---	--

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

4.1.6.2.4 ip/ipv6 host

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

4.1.6.2.5 ip/ipv6 map-hostname

	<pre>{ip ipv6} map-hostname</pre> <pre>no {ip ipv6} map-hostname</pre> Maps between the currently-configured hostname and the loopback address 127.0.0.1. The no form of the command clears the mapping.	
Syntax Description	N/A	
Default	IPv4 mapping is enabled by default IPv6 mapping is disabled by default	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config)# ip map-hostname	
Related Commands	show hosts	

Notes	<ul style="list-style-type: none"> • If no mapping is configured, a mapping between the hostname and the IPv4 loopback address 127.0.0.1 will be added • The no form of the command maps the hostname to the IPv6 loopback address if there is no statically configured mapping from the hostname to an IPv6 address (disabled by default) • Static host mappings are preferred over DNS results. As a result, with this option set, you will not be able to look up your hostname on your configured DNS server; but without it set, some problems may arise if your hostname cannot be looked up in DNS.
-------	---

4.1.6.2.6 show hosts

	<pre>show hosts</pre> <p>Displays hostname, DNS configuration, and static host mappings.</p>
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	<p>3.1.0000</p> <p>3.8.1000: Updated example</p>
Example	<pre>switch (config)# show hosts Hostname: switch1 Name servers: 10.7.77.192 dynamic (DHCP on mgmt0) 10.7.77.135 dynamic (DHCP on mgmt0) 10.198.0.169 dynamic (DHCP on mgmt0) (*) 10.211.0.124 dynamic (DHCP on mgmt0) Domain names: mtl.labs.mlnx dynamic (DHCP on mgmt0) (*) Inactive due to system limits on name servers and domain names. Static IPv4 host mappings: 10.7.144.133 --> switch1 127.0.0.1 --> localhost Static IPv6 host mappings: ::1 --> localhost6 Automatically map hostname to loopback address : yes Automatically map hostname to IPv6 loopback address: no</pre>
Related Commands	
Notes	

4.1.6.3 Routing

4.1.6.3.1 IP route

	<pre>{ip ipv6} route {<network-prefix> <netmask> <network-prefix>/<masklen>} <next-hop> no ip route {<network-prefix> <netmask> <network-prefix>/<masklen>} <next- hop> Sets a static route for a given IP. The no form of the command deletes the static route.</pre>
--	--

Syntax Description	network-prefix	IPv4 or IPv6 network prefix
	netmask	IPv4 netmask formats are: <ul style="list-style-type: none"> • /24 • 255.255.255.0 IPv6 netmask format is: <ul style="list-style-type: none"> • /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	
Example	<code>switch (config)# ip route 20.20.20.0 255.255.255.0 mgmt0</code>	
Related Commands	show ip route	
Notes		

4.1.6.3.2 ipv6 default-gateway

	<code>ipv6 default-gateway {<ip-address> <ifname>}</code> <code>no ipv6 default-gateway</code> Sets a static default gateway. The no form of the command deletes the default gateway.	
Syntax Description	ip address	The default gateway IP address (IPv6)
	ifname	The interface name (e.g., mgmt0, mgmt1)
Default	N/A	
Configuration Mode	config	
History	3.1.0000 3.2.0500: Removed IPv4 configuration option	
Example	<code>switch (config)# ipv6 default-gateway ::1</code>	
Related Commands	show ip/ipv6 route show ipv6 default-gateway	
Notes	<ul style="list-style-type: none"> • The configured default gateway will not be used if DHCP is enabled • In order to configure ipv4 default-gateway use 'ip route' command. 	

4.1.6.3.3 show ip/ipv6 route

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

Example	
<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</pre>	
Related Commands	ip route
Notes	

4.1.6.3.4 show ipv6 default-gateway

	show ipv6 default-gateway [static] Displays the default gateway.	
Syntax Description	static	Displays the static configuration of the default gateway
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example	<pre>switch (config)# show ipv6 default-gateway Active default gateways: 172.30.0.1 (interface: mgmt0) switch (config)# show ipv6 default-gateway static Configured default gateway: 10.10.10.10</pre>	
Related Commands	ipv6 default-gateway	
Notes	The configured IPv4 default gateway will not be used if DHCP is enable	

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.

4.1.6.4.1 ipv6 neighbor

	ipv6 neighbor <ipv6-address> <ifname> <mac-address> no ipv6 neighbor <ipv6-address> <ifname> <mac-address> Adds a static neighbor entry. The no form of the command deletes the static entry.	
Syntax Description	ipv6-address	The IPv6 address
	ifname	The management interface (i.e. mgmt0, mgmt1)

	mac-address	The MAC address
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config)# ipv6 neighbor 2001:db8:701f::8f9 mgmt0 00:11:22:33:44:55	
Related Commands	show ipv6 neighbor ipv6 route arp clear ipv6 neighbors	
Notes	<ul style="list-style-type: none"> • ARP is used only with IPv4. In IPv6 networks, Neighbor Discovery Protocol (NDP) is used similarly. • Use The no form of the command to remove static entries. Dynamic entries can be cleared via the “clear ipv6 neighbors” command. 	

4.1.6.4.2 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 3.6.4110: Updated command	
Example	switch (config)# clear ipv6 neighbors	
Related Commands	ipv6 neighbor show ipv6 neighbor arp	
Notes	<ul style="list-style-type: none"> • Clearing Neighbor Discovery Protocol (NDP) cache removes only the dynamic entries learned and not the static entries configured • Use the no form of the command to remove static entries 	

4.1.6.4.3 show ipv6 neighbors

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

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

4.1.6.5 DHCP

4.1.6.5.1 ip dhcp

	<pre>ip dhcp {default-gateway yield-to-static hostname <hostname> primary-intf <ifname> send-hostname} no ip dhcp {default-gateway yield-to-static hostname primary-intf send- hostname} Sets global DHCP configuration. The no form of the command deletes the DHCP configuration.</pre>	
Syntax Description	yield-to-static	Does not allow you to install a default gateway from DHCP if there is already a statically configured one.
	hostname	Specifies the hostname to be sent during DHCP client negotiation if send-hostname is enabled.
	primary-intf <ifname>	Sets the interface from which a non-interface-specific configuration (resolver and routes) will be accepted via DHCP. Default: "primary-intf mgmt0"
	send-hostname	Enables the DHCP client to send a hostname during negotiation.
Default	Disabled	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config)# ip dhcp default-gateway yield-to-static	
Related Commands	<pre>show ip dhcp dhcp [renew]</pre>	
Notes	DHCP is supported for IPv4 networks only	

4.1.6.5.2 show ip dhcp

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

Example	<pre>switch (config)# show ip dhcp ----- Interface DHCP DHCP Valid Enabled Running lease ----- dummy0 no no no lo no no no mgmt0 yes yes yes mgmt1 no no no mgmts0 no no no mgmts1 no no no vif1 no no no IPv4 dhcp default gateway yields to static configuration: no DHCP primary interface: Configured: mgmt0 Active: mgmt0 DHCP client options: Send Hostname: no Client Hostname: 1.1.1.1</pre>
Related Commands	<pre>ip dhcp dhcp [renew]</pre>
Notes	

4.1.6.6 IP Diagnostic Tools

4.1.6.6.1 ping

	<pre>ping [-LRUbdfnqrVvA] [-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.</pre>	
Syntax Description	Linux Ping options	https://www.lifewire.com/uses-of-command-ping-2201076
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	<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 ... --- 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</pre>	
Related Commands	tracert	
Notes		

4.1.6.6.2 traceroute

	traceroute [-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	-4	Uses IPv4
	-6	Uses IPv6
	-d	Enables socket level debugging
	-F	Sets DF (do not fragment bit) on
	-I	Uses ICMP ECHO for tracerouting
	-T	Uses TCP SYN for tracerouting
	-U	Uses UDP datagram (default) for tracerouting
	-n	Does not resolve IP addresses to their domain names
	-r	Bypasses the normal routing and send directly to a host on an attached network
	-A	Performs AS path lookups in routing registries and print results directly after the corresponding addresses
	-V	Prints version info and exit
	-f	Starts from the first_ttl hop (instead from 1)
	-g	Routes packets through the specified gateway (maximum 8 for IPv4 and 127 for IPv6)
	-i	Specifies a network interface with which to operate
	-m	Sets the max number of hops (max TTL to be reached). Default is 30.
	-N	Sets the number of probes to be tried simultaneously (default is 16)
	-p	Uses destination port. It is an initial value for the UDP destination port (incremented by each probe, default is 33434), for the ICMP seq number (incremented as well, default from 1), and the constant destination port for TCP tries (default is 80).
	-t	Sets the TOS (IPv4 type of service) or TC (IPv6 traffic class) value for outgoing packets
	-l	Uses specified flow_label for IPv6 packets
	-w	Sets the number of seconds to wait for response to a probe (default is 5.0). Non-integer (float point) values allowed too.
-s	Uses source src_addr for outgoing packets.	
-q	Sets the number of probes per each hop. Default is 3.	
-z	Sets minimal time interval between probes (default is 0). If the value is more than 10, then it specifies a number in milliseconds, else it is a number of seconds (float point values allowed too).	
Default	N/A	
Configuration Mode	config	
History	3.1.0000	

Example	
<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 ms6 192.168.10.70 (192.168.10.70) 16.377 ms 16.091 ms 20.475 ms</pre>	
Related Commands	ping
Notes	

4.1.6.6.3 tcpdump

	<p>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]</p> <p>Invokes standard binary, passing command line parameters straight through. Runs in foreground, printing packets as they arrive, until the user hits Ctrl+C.</p>
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
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</pre>
Related Commands	
Notes	

4.1.7 Control Plane Policing (CoPP)



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

This software implements several CoPP mechanisms:

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

- IP filter tables limit the traffic to the CPU coming in from the management ports

4.1.7.1 IP Table Filtering

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

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

Each IP table rule is defined by key, priority, and action:

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

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

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

IP table rules get a lower priority than ACL mechanism.

In the rare case that IP filter is used while the input policy is "drop" (i.e., ip filter chain input policy drop) and an NTP server or an InfiniBand switch with SM HA enabled is used, then the following rule needs to be added that allows src-ip 127.0.0.1 (which is a requirement for any clustered application (e.g., sm-ha) and NTP):

```
ip filter chain input rule append tail target accept dup-delete source-addr 127.0.0.1 /32
```

4.1.7.1.1 Configuring IP Table Filtering

Prerequisite for IPv6:

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

To configure IPv4 table filtering:

1. Select the policy that applies to the input/output chain (default is "accept").

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

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

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

3. Enable IP table filtering.

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

4. Verify IP table filtering configuration.

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

Packet filtering for IPv4: enabled

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

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

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

4.1.7.1.2 Modifying IP Table Filtering

To modify IP table filtering configuration:

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

To delete an existing IP table filtering rule:

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

To delete all existing IP table filtering rules:

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

To insert an IP table filtering rule in a chain:

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

4.1.7.1.3 Rate-Limit Rule Configuration

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

Rate limits can be set using the parameter "rate-limit-above" in order to drop packets whenever traffic is above the set limit. For example: ip filter chain input rule append tail target drop rate-limit-above 1/second source-addr 1.1.1.1 /32.

Another option is to use the parameter "rate-limit". This should be followed by a rule that drops additional packets of the same "type". Alternatively, this can be implicitly achieved by setting the chain policy to "drop" so that it drops packets not processed by matching rules. Otherwise, no effect of the rule is observed as the remaining traffic simply gets accepted.

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

4.1.7.1.4 IP Table Filtering Default Rules

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

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

4.1.7.1.4.1 Firewall Default Rules

Prerouting-Mangle Chain Rules
<ul style="list-style-type: none">• ip filter chain prerouting-mangle rule append tail target drop in-intf mgmt0 protocol tcp conntrack new tcp-op-mss mss-not-in-range 536:65535 not-dest-port 22-23
Input Chain Rules

- ip filter chain input rule append tail target accept in-intf lo
- ip filter chain input rule append tail target drop dup-delete dest-addr 127.0.0.0 /8 in-intf mgmt0
- ip filter chain input rule append tail target accept dup-delete in-intf mgmt0 state established,related
- ip filter chain input rule append tail target drop dup-delete not-dest-port 22-23 in-intf mgmt0 protocol tcp state new tcp-op syn match-not-syn
- ip filter chain input rule append tail target drop dup-delete in-intf mgmt0 fragment enable
- ip filter chain input rule append tail target drop dup-delete in-intf mgmt0 protocol tcp tcp-op flags all
- ip filter chain input rule append tail target drop dup-delete in-intf mgmt0 protocol tcp tcp-op flags none
- ip filter chain input rule append tail target drop dup-delete in-intf mgmt0 state invalid
- ip filter chain input rule append tail target drop dup-delete in-intf mgmt0 protocol tcp tcp-op flags reset rate-limit-above 2/second burst-limit-above 2
- ip filter chain input rule append tail target drop dup-delete in-intf mgmt0 protocol tcp state new rate-limit-above 50/second burst-limit-above 50
- ip filter chain input rule append tail target drop dup-delete in-intf mgmt0 protocol tcp conntrack new rate-limit-above 60/second burst-limit-above 20
- ip filter chain input rule append tail target drop dup-delete in-intf mgmt0 recent name portscan recent rcheck-sec 86400
- ip filter chain input rule append tail target none dup-delete in-intf mgmt0 recent name portscan recent remove
- ip filter chain input rule append tail target none dup-delete dest-port 22 in-intf mgmt0 protocol tcp conntrack new recent set
- ip filter chain input rule append tail target drop dup-delete dest-port 22 in-intf mgmt0 protocol tcp conntrack new recent update-sec 60 recent hitcount 50
- ip filter chain input rule append tail target none dup-delete dest-port 23 in-intf mgmt0 protocol tcp conntrack new recent set
- ip filter chain input rule append tail target drop dup-delete dest-port 23 in-intf mgmt0 protocol tcp conntrack new recent update-sec 60 recent hitcount 50
- ip filter chain input rule append tail target none dup-delete dest-port 443 in-intf mgmt0 protocol tcp conntrack new recent set
- ip filter chain input rule append tail target drop dup-delete dest-port 443 in-intf mgmt0 protocol tcp conntrack new recent update-sec 60 recent hitcount 150
- ip filter chain input rule append tail target none dup-delete dest-port 80 in-intf mgmt0 protocol tcp conntrack new recent set
- ip filter chain input rule append tail target drop dup-delete dest-port 80 in-intf mgmt0 protocol tcp conntrack new recent update-sec 60 recent hitcount 150
- ip filter chain input rule append tail target none dup-delete dest-port 23108 in-intf mgmt0 protocol tcp conntrack new recent set
- ip filter chain input rule append tail target drop dup-delete dest-port 23108 in-intf mgmt0 protocol tcp conntrack new recent update-sec 60 recent hitcount 150
- ip filter chain input rule append tail target none dup-delete dest-port 161 in-intf mgmt0 protocol udp conntrack new recent set
- ip filter chain input rule append tail target drop dup-delete dest-port 161 in-intf mgmt0 protocol udp conntrack new recent update-sec 60 recent hitcount 100
- ip filter chain input rule append tail target accept dup-delete dest-port 22 in-intf mgmt0 protocol tcp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 23 in-intf mgmt0 protocol tcp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 443 in-intf mgmt0 protocol tcp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 80 in-intf mgmt0 protocol tcp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 23108 in-intf mgmt0 protocol tcp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 179 in-intf mgmt0 protocol tcp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 68 in-intf mgmt0 protocol udp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 122 in-intf mgmt0 protocol udp conntrack new,established

- ip filter chain input rule append tail target accept dup-delete dest-port 161 in-intf mgmt0 protocol udp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 6306 in-intf mgmt0 protocol udp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 69 in-intf mgmt0 protocol udp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 389 in-intf mgmt0 protocol udp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 389 in-intf mgmt0 protocol tcp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 1812-1813 in-intf mgmt0 protocol udp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 49 in-intf mgmt0 protocol udp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 49 in-intf mgmt0 protocol tcp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete source-port 53 in-intf mgmt0 protocol udp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete source-port 53 in-intf mgmt0 protocol tcp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 500 in-intf mgmt0 protocol udp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 4500 in-intf mgmt0 protocol udp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 1293 in-intf mgmt0 protocol udp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 1293 in-intf mgmt0 protocol tcp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 1707 in-intf mgmt0 protocol udp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 1707 in-intf mgmt0 protocol tcp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 3786 in-intf lo protocol udp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 33000 in-intf lo protocol udp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete in-intf mgmt0 protocol icmp
- ip filter chain input rule append tail target accept dup-delete source-port 5353 dest-port 5353 in-intf mgmt0 protocol udp conntrack new,established
- ip filter chain input rule append tail target reject-with icmp-port-unreachable dup-delete dest-port 33434-33523 in-intf mgmt0 protocol udp
- ip filter chain input rule append tail target accept dup-delete dest-port 123 in-intf mgmt0 protocol udp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 514 in-intf mgmt0 protocol udp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete dest-port 67 in-intf mgmt0 protocol udp conntrack new,established
- ip filter chain input rule append tail target accept dup-delete comment "Feature HA port" dest-port 60102 in-intf mgmt0 protocol tcp conntrack new,established
- ip filter chain input rule append tail target accept dest-port 636 in-intf mgmt0 protocol udp conntrack new,established
- ip filter chain input rule append tail target accept dest-port 636 in-intf mgmt0 protocol tcp conntrack new,established
- ip filter chain input rule append tail target logging dup-delete in-intf mgmt0

Output Chain Rules

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

Logging Chain Rules

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

4.1.7.2 Control Plane Policing Commands

4.1.7.2.1 ip filter enable | ipv6 filter enable

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

4.1.7.2.2 ip filter chain policy | ipv6 filter chain policy

	{ip ipv6} filter chain <chain_name> policy {accept drop} no {ip ipv6} filter chain <chain_name> policy Configures default policy for a specific chain (if no rule matches this default policy action shall apply). The no form of the command resets default policy for a specific chain.	
Syntax Description	chain_name	Selects a chain for which to add or modify a filter: <ul style="list-style-type: none"> • input - input chain or ingress interfaces • output - output chain or egress interfaces
	accept	Accepts all traffic by default for this chain
	drop	Drops all traffic by default for this chain
Default	Accept for input and output chains	
Configuration Mode	config	
History	3.5.1000	
Example	switch (config) # ipv6 filter chain input policy accept	
Related Commands		
Notes		

4.1.7.2.3 ip filter chain rule target | ipv6 filter chain rule target

	<p>no {ip ipv6} filter chain <chain_name> rule {<number> all} Inserts rule before specified rule number. The no form of the command deletes rule for a specific chain.</p>	
Syntax Description	chain_name	<p>A chain to which to add or modify a filter:</p> <ul style="list-style-type: none"> input - input chain or ingress interfaces output - output chain or egress interfaces
	rule	<ul style="list-style-type: none"> append tail - appends operation to the bottom of operation list insert <oper_num> - inserts operation at specified position (existing operation at that position moves back in the list) modify <oper_num> - modifies existing operation at specified position. Only the parameters specified in this invocation are altered; everything else is left untouched. move <oper_num1> to <oper_num2> - moves one operation to another place in the operation list set <oper_num> - sets operation at specified position (overwrites existing)
	target	<ul style="list-style-type: none"> accept - allows the packets that match the rule into the management plane drop - drops packets that match the rule rate-limit - allows with rate limiting in packets per sec (PPS) reject-with - drops the packet and replies with an ICMP error message
	param	<ul style="list-style-type: none"> rate-limit /[second minute hour] - matches is traffic is less than the set limit rate-limit-above /[second minute hour] - matches is traffic is more than the set limit burst-limit - Maximum initial number of packets to match when setting rate-limit. The default is 5. burst-limit-above - Maximum initial number of packets to match when setting rate-limit-above. The default is 5. comment <text> - specifies description string for this rule (60 chars max) dest-addr <ip> - IP matching a specific destination address or address range. A specific IPv4 address can be provided or an entire subnet by giving an address along with netmask in dot notation or as a CIDR notation (e.g. /24). not-dest-addr <ip> - IP not matching a specific destination address range dest-port <port(s)> - matching a specific destination port or port range not-dest-port <port(s)> - port not matching a specific destination port or port range dup-delete - deletes any preexisting duplicates of this rule in-intf - interface matching a specific inbound interface not-in-intf <if_name> - interface not matching a specific inbound interface out-intf <if_name> - matches a specific outbound interface not-out-intf <if_name> - interface not matching a specific outbound interface

	<p>param4 (cont.)</p> <ul style="list-style-type: none"> • protocol <if_name> - matches a specific protocol • tcp • udp • icmp • all • not-protocol <protocol> - does not match a specific protocol • tcp • udp • icmp • all • source-addr <ip> - matches a specific source address range • not-source-addr <ip> - does not match a specific source address range • source-port <port(s)> - matches a specific source port or port range • not-source-port <port(s)> - does not match a specific source port or port range • state - matches packets in a particular state. Possible values: • established - packet associated with an established connection which has seen traffic in both directions • related - packet that starts a new connection but is related to an existing connection • new - packet that starts a new, unrelated connection • A combination can be entered separated by commas
Default	N/A
Configuration Mode	config
History	3.5.1000
Example	switch (config) # ipv6 filter enable chain input rule append tail target drop state related protocol all dup-delete
Related Commands	
Notes	<ul style="list-style-type: none"> • The source and destination ports may each be either a single number, or a range specified as “<low>-<high>”. For example: “10-20” would specify ports 10 through 20 (inclusive). • The port parameter only works in conjunction with TCP and UDP • Setting a “positive” rule removes any corresponding “not-” rules, and vice-versa • The “state” parameter is a classification of the packet relative to existing connections • If TCP or UDP are selected for the “protocol” parameter, source and/or destination ports may be specified. If ICMP is selected, these options are either ignored, or an error is produced.

4.1.7.2.4 ip filter options include-bridges

	<pre>{ip ipv6} filter options include-bridges no {ip ipv6} filter options include-bridges Applies IP filters to bridges</pre>
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.5.1000
Example	switch (config) # ip filter options include-bridges

Related Commands	
Notes	

4.1.7.2.5 ip filter reset-to-default-rules

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

4.1.7.2.6 show ip filter

	show ip filter Displays IPv4 filtering state.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000
Example	<pre>switch (config) # show ip filter Packet filtering for IPv4: enabled Active IPv4 filtering rules (omitting any not from configuration): Chain 'input' Policy 'accept': Rule 1: Target : accept Protocol : all Source : all Destination : 1.1.1.0/24 Interface : all State : any Other Filter: - Chain 'output' Policy 'accept': Rule 1: Target : reject-with icmp-net-unreachable Protocol : tcp Source : all Destination : all Interface : all State : any Other Filter: dest-port 1000</pre>
Related Commands	
Notes	

4.1.7.2.7 show ip filter all

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

4.1.7.2.8 show ip filter configured

	show ip filter configured Displays IPv4 filtering configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000
Example	<pre>switch (config) # show ip filter configured Packet filtering for IPv4: enabled IPv4 configuration: Chain 'input' Policy 'accept': Rule 1: Target : accept Protocol : all Source : all Destination : 1.1.1.0/24 Interface : all State : any Other Filter: - Chain 'output' Policy 'accept': Rule 1: Target : reject-with icmp-net-unreachable Protocol : tcp Source : all Destination : all Interface : all State : any Other Filter: dest-port 1000</pre>

Related Commands	
Notes	

4.1.7.2.9 show ipv6 filter

	show ipv6 filter Displays IPv6 filtering state.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000
Example	<pre>switch (config) # show ipv6 filter Packet filtering for IPv6: enables Active IPv6 filtering rules (omitting any not from configuration): Chain 'input' Policy 'accept': Rule 1: Target : accept Protocol : all Source : all Destination : 1.1.1.0/24 Interface : all State : any Other Filter: - Chain 'output' Policy 'accept': Rule 1: Target : reject-with icmp-net-unreachable Protocol : tcp Source : all Destination : all Interface : all State : any Other Filter: dest-port 1000</pre>
Related Commands	
Notes	

4.1.7.2.10 show ipv6 filter all

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

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

4.1.7.2.11 show ipv6 filter configured

	<pre>show ipv6 filter configured Displays IPv6 filtering configuration.</pre>
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000
Example	<pre>switch (config) # show ipv6 filter configured Packet filtering for IPv6: enables IPv6 configuration: Chain 'input' Policy 'accept': Rule 1: Target : accept Protocol : all Source : all Destination : 1.1.1.0/24 Interface : all State : any Other Filter: - Chain 'output' Policy 'accept': Rule 1: Target : reject-with icmp-net-unreachable Protocol : tcp Source : all Destination : all Interface : all State : any Other Filter: dest-port 1000</pre>
Related Commands	
Notes	

4.2 Chassis Management



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

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

Additionally, it monitors:

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

4.2.1 System Health Monitor

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

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

4.2.1.1 Re-Notification on Errors

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

4.2.1.2 System Health Monitor Alerts Scenarios

System Health Monitor sends notification alerts in the following cases:

Alert Message	Scenario	Notification Indicator	Recovery Action	Recovery Message
Mismatch detected in the airflow direction of a Fan or PS Fan module	Fan or PS fan has a different airflow direction than other fans	Email, System status LED set to red, log alert, SNMP	Change the fan module with another fan module with same airflow direction to align with other fans and ps fans	"System airflow direction is aligned"
<fan_name> speed is below minimal range	A chassis fan speed is below minimal threshold: 15% of maximum speed	Email, fan LED and system status LED set red, log alert, SNMP.	Check the fan and replace it if required	"<fan_name> has been restored to its normal state"
Fan <fan_number> speed in spine number <spine_number> is below minimal range	A spine fan speed is below minimal threshold: 30% of maximum speed	Email, fan LED and system status LED set red, log alert, SNMP	Check the fan and replace it if required	"Fan speed <fan_number> in spine number <spine_number> has been restored to its normal state"
<fan_name> is unresponsive	A chassis fan is not responsive on the switch system	Email, fan LED and system status LED set red, log alert, SNMP	Check fan connectivity and replace it if required	"<fan_name> has been restored to its normal state"
Fan <fan_number> in spine number <spine_number> is unresponsive	A spine fan is not responsive on the switch system	Email, fan LED and system status LED set red, log alert, SNMP	Check fan connectivity and replace it if required	"Fan <fan_number> in spine number <spine_number> has been restored to its normal state"
<fan_name> is not present	A chassis fan is missing	Email, fan LED and system status LED set red, log alert, SNMP	Insert a fan unit	"<fan_name> has been restored to its normal state"
Fan <fan_number> in spine number <spine_number> is not present.	A spine fan is missing	Email, fan LED and system status LED set red, log alert, SNMP	Insert a fan unit	"Fan <fan_number> in spine number <spine_number> has been restored to its normal state"
Insufficient number of working fans in the system	Insufficient number of working fans in the system	Email, fan LED and system status LED set red, log alert, SNMP	Plug in additional fans or change faulty fans	"The system currently has sufficient number of working fans"
Power Supply <ps_number> voltage is out of range	The power supply voltage is out of range.	Email, power supply LED and system status LED set red, log alert, SNMP	Check the power connection of the PS	"Power Supply <ps_number> voltage is in range"
Power supply <ps_number> temperature is too hot	A power supply unit temperature is higher than the maximum threshold of 70 Celsius on the switch system	Email, power supply LED and system status LED set red, log alert, SNMP	Check chassis fans connections. On switch systems, check system fan connections.	"Power supply <ps_number> temperature is back to normal"

Alert Message	Scenario	Notification Indicator	Recovery Action	Recovery Message
Power Supply <number> is unresponsive	A power supply is malfunctioning or disconnected	Email, system status and power supply LED set red, log alert, SNMP	Connect power cable or replace malfunctioning PS	“Power supply has been removed” or “PS has been restored to its normal state”
Unit/leaf/spine <leaf/spine number> is unresponsive	A leaf/spine is not responsive	Email, system status LED set red, log alert, SNMP	Check leaf/spine connectivity and replace it if required	“Leaf/spine number <leaf/spine number> has been restored to its normal state”
Unit/leaf/spine voltage is out of range	One of the voltages on the switch system is below minimal threshold or higher than the maximum threshold - both thresholds are 15% of the expected voltage	Email, system status LED set red, log alert, SNMP	Check leaf connectivity	“Unit voltage is in range”
ASIC temperature is too hot	An ASIC unit temperature is higher than the maximum threshold of 105 Celsius on switch systems	Email, system status LED set red, log alert, SNMP	Check the fan’s system	“ASIC temperature is back to normal”

4.2.2 Power Management

4.2.2.1 Power Supply Options

MLNX-OS offers power redundancy configurations and monitoring for modular switch systems. Modular switch systems have the following redundancy configuration modes:

- “combined”—no power supply is reserved. The redundancy is not enabled.
- “ps-redundant”—one power supply unit is redundant to the rest. The system can work with one less power supply unit.
- “grid-redundant”—the power supplies are split into two logical power supply grids, first half of the PSUs belongs to grid A and the second half to grid B. The systems can work with only one grid. When using grid-redundancy mode the power budget is calculated according to the minimum power budget between the grids. This mode is available on CS75xx chassis systems. During switch initialization, or hot-plugging of switch components, MLNX-OS enables and/or disables switch components according to the available power budget.

MLNX-OS may send power alarms (via SNMP or email) as follow:

- If the available budget is insufficient for all the system components an “insufficientPower” event is generated. In this mode several switch components may be disabled.

- If the total power of the system is insufficient for redundancy, a “lowPower” event is generated
- If a connected power supply provides below 1.6K Watts or grid-redundancy mode is configured and a power supply is connected to a 110V grid, then a “powerRedundancyMismatch” event is generated, where grid redundancy can not be achieved in such configuration.

In case of an insufficient-power mode, the order in which the FRUs are turned ON is first spines (1,2,3...max) and then the leafs (1,2,3...max), while the order of the FRUs in case of turning them OFF is first the spines (max...3) and then the leafs (max...1). The management modules are not affected.

For the trap OID, please refer to the Mellanox-MIB file.

Power cycle is needed after changing power redundancy mode on a modular switch system.

4.2.2.2 Width Reduction Power Saving

Link width reduction (LWR) is a NVIDIA

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

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

LWR is relevant only for 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.

The following table describes LWR configuration behavior:

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

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

4.2.2.3 Managing Chassis Power

It is possible to power down or power up modules in a chassis by using the commands “power enable” and “no power enable”.

1. Run the command “show power” to get a list of modules that are available to power up or down.
2. To power down a desired module, run:

```
switch (config) # no power enable <module>
```

3. To power up a desired module, run:

```
switch (config) # power enable <module>
```

4. Using the “show power” command it is possible to see the power consumption of the system and also the power consumption by power supply unit.

4.2.3 Monitoring Environmental Conditions

1. Display module’s temperature.

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

2. Display measured voltage levels of power supplies.

```
switch (config) # show voltage
-----
Module  Power Meter      Reg                Expected  Actual  Status  High  Low
              Voltage          Voltage
-----
MGMT    acdc-monitor1    DDR3 0.675V        0.68     0.67    OK       0.78  0.57
MGMT    acdc-monitor1    CPU 0.9V           0.90     0.86    OK       1.03  0.77
MGMT    acdc-monitor1    SYS 3.3V           3.30     3.36    OK       3.79  2.80
```

```

MGMT acdc-monitor1 CPU 1.8V 1.80 1.82 OK 2.07 1.53
MGMT acdc-monitor1 CPU/PCH 1.05V 1.05 1.06 OK 1.21 0.89
MGMT acdc-monitor1 CPU 1.05V 1.05 1.06 OK 1.21 0.89
MGMT acdc-monitor1 DDR3 1.35V 1.35 1.35 OK 1.55 1.15
MGMT acdc-monitor1 USB 5V 5.00 5.04 OK 5.75 4.25
MGMT acdc-monitor1 1.05V LAN 1.50 1.51 OK 1.72 1.27
MGMT ASICVoltMonitor1 Asic 1.2V 1.20 1.21 OK 1.38 1.02
MGMT ASICVoltMonitor1 Asic 3.3V 3.30 3.31 OK 3.79 2.80
MGMT ASICVoltMonitor2 Vcore SX 0.95 0.96 OK 1.09 0.81
MGMT ASICVoltMonitor2 Asic 1.8V 1.80 1.81 OK 2.07 1.53
MGMT acdc-monitor2 3.3V Switch IB 3.30 3.36 OK 3.79 2.80
PS1 power-mon vout 12V 12.00 12.07 OK 13.80 10.20

```

3. Display the fan speed and status. Run:

```

switch (config) # show fan
-----
Module      Device      Fan  Speed (RPM)  Status
-----
FAN1        FAN         F1   6297.00      OK
FAN1        FAN         F2   5421.00      OK
FAN2        FAN         F1   6355.00      OK
FAN2        FAN         F2   5378.00      OK
FAN3        FAN         F1   6183.00      OK
FAN3        FAN         F2   5421.00      OK
FAN4        FAN         F1   6268.00      OK
FAN4        FAN         F2   5399.00      OK
PS1         FAN         F1   10336.00     OK
PS2         FAN         -    -            NOT PRESENT

```

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

```

switch (config) # show power consumers
-----
Module Device      Sensor  Power [Watts]  Voltage [Volts]  Current [Amp]  Status
-----
PS1   power-mon  input  39.94    12.07    3.31      OK
MGMT  acdc-monitor1  input  2.11    12.00    0.18      OK

Total power used : 42.05 Watts

```

4.2.4 USB Access

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

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

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

Examples:

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

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

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

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

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

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

4.2.5 Unit Identification LED

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

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

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

To verify the LED status, run:

```
switch (config) # show leds
Module  LED           Status
-----
MGMT    STATUS         Green
MGMT    FAN1           Green
MGMT    FAN2           Green
MGMT    FAN3           Green
MGMT    FAN4           Green
MGMT    PS_STATUS      Green
MGMT    PS1            Green
MGMT    PS2            Green
MGMT    UID            Blue
```

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

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

4.2.6 High Availability (HA)

NVIDIA high end management modular switch systems support redundant management modules. Chassis HA reduces downtime as it assures continuity of the work even when a management module dies. Chassis HA management allows the systems administrator to associate a single IP address with the appliance. Connecting to that IP address allows the user to change and review the system's chassis parameters regardless of the active management module.

4.2.6.1 Chassis High Availability Nodes Roles

Every node in the Chassis HA has one of the following roles/modes:

- Master—the node that manages chassis configurations and services the chassis IP addresses
- Slave—the node that replaces the Master node and takes over its responsibilities once the Master node is down

The master node is the only node that has access to chassis components such as temperature, inventory and firmware.

The CPU role of the current management node can be recognized by following one these methods:

- Running the command “show chassis ha”

```
switch (config) # show chassis ha
2-node HA state:
Box management IPv4: 10.7.146.44/24
Box management IPv6: fdfd:fdfd:7:145::1033:47fd/64
interface       : mgmt0
local role      : master
local slot      : 1
other state     : not-present
reset count     : 0
```

- Check the LEDs in the management modules as displayed in the figure below
- Go to the WebUI → System → Modules page and see the information on the LEDs

4.2.6.2 Malfunctioned CPU Behavior

When a CPU is not responding to an internal communication with the other CPU, the non responding CPU will be reset by the other CPU. Each time a CPU resets, a counter is incremented. After 5 resets a CPU is considered malfunctioned and will be shut down.

To verify how many times a CPU is reset, run:

```
switch [default: master] (config) # show chassis ha
2-node HA state:
Box management IPv4: 10.7.146.44/24
Box management IPv6: fdfd:fdfd:7:145::1033:47fd/64
interface       : mgmt0
local role      : master
local slot      : 1
other state     : not-present
reset count     : 0
```

To verify if a CPU has been shut down, either run:

```
switch [default: master] (config) # show chassis ha
2-node HA state:
Box management IPv4: 10.7.146.44/24
Box management IPv6: fdfd:fdfd:7:145::1033:47fd/64
interface       : mgmt0
local role      : master
local slot      : 1
other state     : not-present
reset count     : 0
```

Or check the system page in the WebUI, the management figure will be grayed out. To enable the malfunctioned CPU, first replace it and run “chassis ha reset other”.

4.2.6.3 Box IP Centralized Location

Box IP (BIP) centralized management infrastructure enables you to configure and monitor the system. The BIP continues to function even if one of the management blades dies. Box IP is defined by running the command “chassis ha bip <board IP address>”. The created BIP is used as the master IP’s alias. For example:

```
switch [standalone: master] (config) # chassis ha bip 192.168.10.100 255.255.255.0
```

4.2.6.4 System Configuration

System configuration changes should be performed by the master using the BIP otherwise they are overridden by the master configuration.

Chassis HA is based on database replication enabling the entire master configuration to be replicated to the slave. Data such as chassis configuration is replicated. However, run time

information such as time, logs, active user lists, is not copied. Additionally, node specific configuration information such as host name and IP address is not copied.

Chassis HA requires connectivity of both management modules (mgmt0, mgmt1) in the same broadcast domain.

The SM commands are only visible to the SM HA master in a modular system. This is node would display "master" in its CLI prompt.

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

If the node shows "slave" or "unknown", the node is not the "master" and thus would not be able to use the IB SM commands.

"unknown" indicates that mgmt0 is not LinkUp and is not assigned a valid IPv4 address. On modular systems, the mgmt0 interface on all installed management modules must be:

- LinkUp
- With a valid IPv4 address
- In the same L2 broadcast domain

Even if only one module is installed, it must have a mgmt0 interface that is LinkUp and with a valid IPv4 address.

4.2.6.5 Takeover Functionally

Management CPU functional takeover takes up to 20-30 seconds. However, when plugging in a module, you need to wait for approximately 3 minutes before making any other hardware change. During the takeover process, the Master LED status is differentiated by a color scheme. To verify the system's status, run the "show chassis ha" command on both managements.

If the CPU malfunctions, the system resets it 5 times in an attempt to solve the issue. If the CPU is not activated after the reset, the system powers it off as well as its attached spine. Once the CPU is powered off, the user should replace the malfunctioned CPU module. To power on the CPU and the attached spine, plug the module in, log into the Master CPU and run the "chassis ha power enable other" command.

Although the LEDs are functional during the takeover, wait for approximately 3 minutes before making any other hardware change.

Master example:

```
switch [default: master] (config) # show chassis ha
2-node HA state:
Box management IPv4: 10.7.146.44/24
Box management IPv6: fdfd:fdfd:7:145::1033:47fd/64
interface       : mgmt0
local role      : master
local slot     : 1
other state     : not-present
reset count    : 0
```

Slave example:

```
switch [default: master] (config) # show chassis ha
2-node HA state:
```

```

Box management IPv4: 10.7.146.44/24
Box management IPv6: fdfd:fdfd:7:145::1033:47fd/64
interface       : mgmt0
local role      : master
local slot      : 1
other state     : not-present
reset count     : 0

```

Not following these instructions may result in some errors in the log. These errors may be safely ignored.

4.2.7 System Reboot

4.2.7.1 Rebooting 1U Switches

To reboot a 1U switch system, run:

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

4.2.7.2 Rebooting Modular Switches

NVIDIA high end management modular switch systems support redundant management modules. Chassis HA reduces downtime as it assures continuity of the work even when a management module dies. Chassis HA management allows the systems administrator to associate a single IP address with the appliance. Connecting to that IP address allows the user to change and review the system's chassis parameters regardless of the active management module.

To reboot modular switches:

1. Connect to the box IP (BIP). Please refer to [“Box IP Centralized Location”](#) for more information.
2. Reboot the slave management. Run:

```
switch [default: master] (config) # chassis ha reset other
```

3. Reboot the master management. Run:

```
switch [default: master] (config) # reload
```

4.2.8 Viewing Active Events

The OS supports viewing all active events on the system. The following events may be observed with the command [“show system hardware events”](#).

Event Name	Description
Ethernet Family	
Invalid Mac (SMAC=MC)	Source MAC is a multicast address

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

4.2.9 Chassis Management Commands



- [4.2.9.1 Chassis Management](#)
 - [4.2.9.1.1 clear counters](#)
 - [4.2.9.1.2 health](#)
 - [4.2.9.1.3 led uid](#)
 - [4.2.9.1.4 power enable](#)
 - [4.2.9.1.5 system manage inband-ib](#)
 - [4.2.9.1.6 power redundancy-mode](#)
 - [4.2.9.1.7 system profile](#)
 - [4.2.9.1.8 usb eject](#)
 - [4.2.9.1.9 show asic-version](#)
 - [4.2.9.1.10 show bios](#)

- [4.2.9.1.11 show cpld](#)
- [4.2.9.1.12 show fan](#)
- [4.2.9.1.13 show health-report](#)
- [4.2.9.1.14 show inventory](#)
- [4.2.9.1.15 show leds](#)
- [4.2.9.1.16 show memory](#)
- [4.2.9.1.17 show module](#)
- [4.2.9.1.18 show power](#)
- [4.2.9.1.19 show power consumers](#)
- [4.2.9.1.20 show protocols](#)
- [4.2.9.1.21 show resources](#)
- [4.2.9.1.22 show system capabilities](#)
- [4.2.9.1.23 show system manage inband-ib](#)
- [4.2.9.1.24 show system profile](#)
- [4.2.9.1.25 show system type](#)
- [4.2.9.1.26 show temperature](#)
- [4.2.9.1.27 show version](#)
- [4.2.9.1.28 show version concise](#)
- [4.2.9.1.29 show voltage](#)
- [4.2.9.2 Chassis High Availability](#)
 - [4.2.9.2.1 chassis ha bip](#)
 - [4.2.9.2.2 chassis ha](#)
 - [4.2.9.2.3 chassis ha power enable other](#)
 - [4.2.9.2.4 show chassis ha](#)
 - [4.2.9.2.5 chassis ha bipv6](#)

4.2.9.1 Chassis Management

4.2.9.1.1 clear counters

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

4.2.9.1.2 health

	health {max-report-len <length> re-notif-cntr <counter> report-clear} Configures health daemon settings.	
Syntax Description	max-report-len <length>	Sets the length of the health report (number of line entries) Range: 10-2048
	re-notif-cntr <counter>	Health control changes notification counter in seconds Range: 120-7200
	report-clear	Clears the health report
Default	max-report-len: 50 re-notif-cntr:	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # health re-notif-cntr 125	
Related Commands	show health-report	
Notes		

4.2.9.1.3 led uid

	led <module> uid <on off> Configures the UID LED.	
Syntax Description	module	Specifies the module whose UID LED to configure
	on	Turns on UID LED
	off	Turns off UID LED
Default	N/A	
Configuration Mode	config	
History	3.6.1002 3.6.2002: Added modular switch support	
Example	switch (config) # led MGMT uid on	
Related Commands		
Notes	<ul style="list-style-type: none"> On 1U switch systems, the module parameter can only be MGMT On modular switch systems, the module parameter may be MGMT#, L#, S# (e.g. MGMT1, L01, S01) 	

4.2.9.1.4 power enable

	power enable <module name> no power enable <module name> Powers on the module. The no form of the command shuts down the module.	
Syntax Description	module name	Enables power for selected module
Default	Power is enabled on all modules	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # power enable L01	
Related Commands	show power show power consumers	
Notes	<ul style="list-style-type: none"> • It is recommended to run this command prior to extracting a module from the switch system, else errors are printed in the log • This command is not applicable on 1U systems 	

4.2.9.1.5 system manage inband-ib

	system manage inband-ib no system manage inband-ib Enables remote inband management of the system. The no form of the command disables remote inband management of the system.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config	
History	3.8.1000	
Example	switch (config) # system manage inband-ib	
Related Commands	show system manage inband-ib	
Notes	This command is available only on Quantum based switch systems	

4.2.9.1.6 power redundancy-mode

	power redundancy-mode [combined grid-redundant ps-redundant] no power redundancy-mode Controls the power supply redundancy mode. The no form of the command resets power redundancy mode to the default value.	
Syntax Description	combined	No redundancy - no alarm threshold
	grid-redundant	N+N - the alarm threshold will be set to a level, indicating when the power availability falls below power that can support N+N scheme

	ps-redundant	N+1 - the alarm threshold will be set to a level, indicating when the power availability falls below power that can support N+1 scheme
Default	N/A	
Configuration Mode	config	
History	3.2.0000 3.10.1000: Added the no form of the command	
Example	switch (config) # power redundancy-mode combined	
Related Commands		
Notes	<ul style="list-style-type: none"> The difference between the modes sets the threshold for power supply redundancy failure. It does not change any power supply configuration. This command is not applicable for 1U or blade systems. 	

4.2.9.1.7 system profile

	system profile {ib-single-switch ib-no-adaptive-routing-single-switch ib [split-ready] [num-of-swids <swid-num>] [no-adaptive-routing] [ib-router] [adaptive-routing-groups <value>]} [force] Sets the profile of the system to InfiniBand with various parameters	
Syntax Description	ib-single-switch	Enables InfiniBand switch profile All network interfaces link protocol set to InfiniBand
	ib-no-adaptive-routing-single-switch	Enables InfiniBand switch profile without adaptive routing capabilities All network interfaces link protocol set to InfiniBand with disabled adaptive routing
	split-ready	Enables the system to reboot in split enable mode with capability to configure 2x the number of ports exposed to IB utilities. Note: This parameter is available only on Quantum-based systems.
	ib-router	Enables IB Routing capability on the system
	num-of-swids	Multiple switch IDs are configurable <ul style="list-style-type: none"> adaptive routing—enables adaptive routing ib-router—enables IB router Note: If num-of-swids is not defined then it is set to 1 by default.
	no-adaptive-routing	Disables adaptive routing
	adaptive-routing-groups	Sets adaptive routing groups. <div style="border: 1px solid orange; padding: 5px; text-align: center;">Allowed only when adaptive routing is enabled.</div>
Default	The default system profile depends on the system.	
Configuration Mode	config	

History	<p>3.1.0000</p> <p>3.2.1100: Added “vpi-single-switch” option</p> <p>3.3.4100: Added SX6036G3.3.4302Added system profile ib-no-adaptive-routing-single-switch</p> <p>3.6.1002: Added system profile “ib num-of-swids”</p> <p>3.6.6162: Added system profile “num of adaptive routing”</p> <p>3.7.0020: Added system profile “ib split-ready” for 1U systems</p> <p>3.8.1100: Updated syntax description for the parameter “adaptive-routing-groups”</p> <p>3.9.0300: Added system profile “ib split-ready” for modular systems</p> <p>3.9.2000: Updated note</p> <p>3.10.6000: Updated note</p>
Example	<code>switch (config) # system profile ib-single-switch</code>
Related Commands	<p>port type</p> <p>show system profile</p> <p>show ports type</p>
Notes	<ul style="list-style-type: none"> • This command requires approval because reboot is performed and all configuration is removed • This command deletes all switch configuration (keeping configuration necessary for network connectivity such as interfaces, routes, and ARP) and resets the system • System profile “ib-no-adaptive-routing-single-switch profile” is the default profile for InfiniBand switches • The parameter “adaptive-routing-groups” is only available when “adaptive-routing” is configured • Refer to the “port type” command in order to change the link protocol • System profile “ib split-ready” must run together with num-of-swids 1 • IB router and adaptive routing are enabled only if specified but cannot be enabled at the same time • IB router only works when adaptive routing is disabled. • In NDR systems, the maximum number of SWIDs is 8.

4.2.9.1.8 usb eject

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

4.2.9.1.9 show asic-version

	<p>show asic-version</p> <p>Displays firmware ASIC version.</p>
Syntax Description	N/A
Default	N/A

Configuration Mode	Any command mode
History	3.1.0000 3.4.2008: Updated example
Example	<pre>switch (config) # show asic-version ===== Module Device Version ===== L05 SIB2-1 15.0200.0092 L05 SIB2-2 15.0200.0092</pre>
Related Commands	
Notes	

4.2.9.1.10 show bios

	show bios Displays the BIOS version information.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.3.4150
Example	<pre>switch (config) # show bios BIOS version : 4.6.5 BIOS subversion : Official AMI Release BIOS release date : 07/02/2021</pre>
Related Commands	
Notes	

4.2.9.1.11 show cpld

	show cpld Displays status of all CPLDs in the system.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.3.4302: Updated example 3.10.1000: Updated example to reflect the part number (PN) field 3.10.1100: Updated example to reflect Version Minor
Example	<pre>switch (config) # show cpld ----- Name Type Version Version Minor PN ----- Cpld1 CPLD_TOR 9 1 0x0078 Cpld2 CPLD_SWB_UNIFIED 3 3 0x0128 Cpld3 CPLD_LED 1 0 0x00d1</pre>
Related Commands	
Notes	

4.2.9.1.12 show fan

	show fan Displays fans status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show fan ----- Module Device Fan Speed Status (RPM) ----- FAN1 FAN F1 6297.00 OK FAN1 FAN F2 5421.00 OK FAN2 FAN F1 6355.00 OK FAN2 FAN F2 5378.00 OK FAN3 FAN F1 6183.00 OK FAN3 FAN F2 5421.00 OK FAN4 FAN F1 6268.00 OK FAN4 FAN F2 5399.00 OK PS1 FAN F1 10336.00 OK PS2 FAN - - NOT PRESENT</pre>
Related Commands	
Notes	

4.2.9.1.13 show health-report

	show health-report Displays health report.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.3.0000: Output update 3.11.2000: Output update
Example	<pre>switch (config) # show health-report ALERTS CONFIGURATION Re-notification counter (sec): [3600] Report max counter : [50] HEALTH REPORT No Health issues file</pre>
Related Commands	health
Notes	

4.2.9.1.14 show inventory

	show inventory Displays system inventory.
--	---

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.4.1604: Removed CPU module output from example 3.5.1000: Removed Type column from example 3.6.1002: Updated example
Example	
<pre>switch (config) # show inventory ----- Module Part Number Serial Number Asic Rev. HW Rev. ----- CHASSIS MSB7800-ES2F MT1602X17464 N/A A1 MGMT MSB7800-ES2F MT1602X17464 0 A1 FAN1 MTEF-FANF-A MT1602X16943 N/A A3 FAN2 MTEF-FANF-A MT1602X16944 N/A A3 FAN3 MTEF-FANF-A MT1602X16956 N/A A3 FAN4 MTEF-FANF-A MT1602X16957 N/A A3 PS1 MTEF-PSF-AC-A MT1601X09908 N/A A3</pre>	
Related Commands	
Notes	

4.2.9.1.15 show leds

	show leds [<module>] Displays the LED status of the switch system.	
Syntax Description	module	Specifies the module whose LED status to display
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.1002 3.6.2002: Updated example	

Example	<pre>switch (config) # show leds Module LED Status ----- MGMT1 STATUS Green MGMT1 REAR_FAN Green MGMT1 PS Green MGMT1 FRONT_FAN Green MGMT1 MASTER/SLAVE Green L01 STATUS Green L01 UID Blue L02 STATUS Green L02 UID Blue L03 STATUS Green L03 UID Off L04 STATUS Green L04 UID Off L05 STATUS Green L05 UID Off L06 STATUS Green L06 UID Off S01 STATUS Green S01 FAN Green S02 STATUS Green S02 FAN Green S03 STATUS Green S03 FAN Green FAN1 STATUS Green FAN2 STATUS Green FAN3 STATUS Green FAN4 STATUS Green</pre>
Related Commands	
Notes	

4.2.9.1.16 show memory

	<pre>show memory Displays memory status.</pre>
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.7.1000: Updated example
Example	
<pre>switch (config) # show memory ----- Memory Space Total Used Free Used+B/C Free-B/C ----- Physical 15848 MB 2849 MB 12999 MB 3854 MB 11994 MB Swap 0 MB 0 MB 0 MB Physical Memory Borrowed for System Buffers and Cache: Buffers : 27 MB Cache : 910 MB Total Buffers/Cache: 937 MB</pre>	
Related Commands	
Notes	

4.2.9.1.17 show module

	show module Displays modules status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.3.0000: Added "Is Fatal" column 3.4.2008: Updated command output 3.4.3000: Updated command output and added note
Example	<pre>switch (config) # show module ===== Module Status ===== MGMT ready FAN1 ready FAN2 ready PS1 ready PS2 not-present</pre>
Related Commands	
Notes	The Status column may have one of the following values: error, fatal, not-present, powered-off, powered-on, ready.

4.2.9.1.18 show power

	show power Displays power supplies and power usage.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.5.1000: Updated example
Example	<pre>switch (config) # show power ----- Module Device Sensor Power Voltage Current Capacity Feed Status [Watts] [Volts] [Amp] [Watts] ----- PS1 power-mon input 32.25 12.11 1.26 800.00 DC OK PS2 power-mon input 46.56 12.13 2.33 800.00 DC OK</pre>
Related Commands	
Notes	

4.2.9.1.19 show power consumers

	show power consumers Displays power consumption information.
--	---

Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.5.1000: Updated example
Example	<pre>switch (config) # show power consumers ----- Module Device Sensor Power Voltage Current Status [Watts] [Volts] [Amp] ----- MGMT CURR_MONITOR 12V 52.96 11.71 4.52 OK PS1 power-mon input 252.00 12.00 20.25 OK PS2 power-mon input 280.00 12.03 23.25 OK Total power used : 52.96 Watts</pre>
Related Commands	
Notes	

4.2.9.1.20 show protocols

	<pre>show protocols Displays all protocols enabled in the system.</pre>
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.2.3000 3.3.4550: Updated example 3.6.1002: Updated example
Example	<pre>switch (config) # show protocols Infiniband enabled sm enabled router disabled</pre>
Related Commands	
Notes	

4.2.9.1.21 show resources

	<pre>show resources Displays system resources.</pre>
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000

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%</pre>
Related Commands	
Notes	

4.2.9.1.22 show system capabilities

	<pre>show system capabilities Displays system capabilities.</pre>
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	<pre>3.1.0000 3.3.0000: Added gateway support 3.6.1002: Updated example 3.7.0000: Updated example</pre>
Example	<pre>switch (config) # show system capabilities IB: Supported, L2, Adaptive Routing, Split Ready Max SM nodes: 648 IB Max licensed speed: EDR</pre>
Related Commands	show system profile
Notes	

4.2.9.1.23 show system manage inband-ib

	<pre>show system manage inband-ib Displays whether inband management over InfiniBand is currently allowed.</pre>
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.8.1000
Example	<pre>switch (config) # show system manage inband-ib Manage inband-ib: no</pre>
Related Commands	system manage inband-ib
Notes	This command is available only on Quantum based switch systems

4.2.9.1.24 show system profile

	<code>show system profile</code> Displays system profile.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.2.0000 3.7.0000: Updated example
Example	<pre>switch (config) # show system profile Profile : ib Number of SWIDs : 1 Adaptive Routing : yes Adaptive Routing Groups : 2048 IB Routing : no</pre>
Related Commands	<code>system profile</code>
Notes	

4.2.9.1.25 show system type

	<code>show system type</code> Displays system type.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.5.1000
Example	<pre>switch (config) # show system type SB7700</pre>
Related Commands	
Notes	

4.2.9.1.26 show temperature

	<code>show temperature</code> Displays system temperature sensors status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000

Example	<pre>switch (config) # show temperature ----- Module Component Reg CurTemp Status (Celsius) ----- MGMT SIB2 T1 32.00 OK MGMT Board AMB temp T1 23.50 OK MGMT Ports AMB temp T1 27.50 OK MGMT CPU package Sensor T1 27.00 OK MGMT CPU Core Sensor T1 18.00 OK MGMT CPU Core Sensor T2 27.00 OK PS1 power-mon T1 22.50 OK</pre>
Related Commands	
Notes	

4.2.9.1.27 show version

	<pre>show version Displays version information for the currently running system image.</pre>
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show version Product name: MLNX-OS Product release: 3.11.1954-007 Build ID: #1-dev Build date: 2023-10-18 15:21:05 Target arch: x86_64 Target hw: x86_64 Version summary: X86_64 3.11.1954-007 2023-10-18 15:21:05 x86_64 Product model: x86onie Host ID: 0C42A117E840 System serial num: MT2006X07803 System UUID: 62cbd568-7d2a-11ea-8000-0c42a1589d10 Uptime: 17h 1m 3.828s CPU load averages: 0.00 / 0.00 / 0.00 Number of CPUs: 4 System memory: 846 MB used / 6954 MB free / 7800 MB total Swap: 0 MB used / 0 MB free / 0 MB total</pre>
Related Commands	
Notes	

4.2.9.1.28 show version concise

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

Example	switch (config) # show version concise X86_64 3.6.4006 2017-07-03 16:17:39 x86_64
Related Commands	
Notes	

4.2.9.1.29 show voltage

	show voltage Displays voltage level measurements on different sensors.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000 3.3.5006: Updated example	
Example	<pre>switch (config) # show voltage ===== Module Power Meter Reg Expected Actual Status High Low Voltage Voltage Voltage Voltage ===== MGMT BOARD_MONITOR USB 5V sensor 5.00 5.15 OK 5.55 4.45 MGMT BOARD_MONITOR Asic I/O sensor 2.27 2.11 OK 2.55 1.99 MGMT BOARD_MONITOR 1.8V sensor 1.80 1.79 OK 2.03 1.57 MGMT BOARD_MONITOR SYS 3.3V sensor 3.30 3.28 OK 3.68 2.92 MGMT BOARD_MONITOR CPU 0.9V sensor 0.90 0.93 OK 1.04 0.76 MGMT BOARD_MONITOR 1.2V sensor 1.20 1.19 OK 1.37 1.03 MGMT CPU_BOARD_MONITOR 12V sensor 12.00 11.67 OK 13.25 10.75 MGMT CPU_BOARD_MONITOR 12V sensor 2.50 2.46 OK 2.80 2.20 MGMT CPU_BOARD_MONITOR 2.5V sensor 3.30 3.26 OK 3.68 2.92 MGMT CPU_BOARD_MONITOR SYS 3.3V sensor 3.30 3.24 OK 3.68 2.92 MGMT CPU_BOARD_MONITOR SYS 3.3V sensor 1.80 1.79 OK 2.03 1.57 MGMT CPU_BOARD_MONITOR 1.8V sensor 1.20 1.24 OK 1.37 1.03</pre>	
Related Commands		
Notes		

4.2.9.2 Chassis High Availability

4.2.9.2.1 chassis ha bip

	chassis ha bip <board-ip-address> Configures Chassis Board IP (BIP).	
Syntax Description	board-ip-address	Sets the chassis virtual IP address
Default	0.0.0.0	
Configuration Mode	config	

History	3.1.0000
Example	<code>switch (config) # chassis ha bip 192.168.10.100</code>
Related Commands	show chassis ha
Notes	This command is applicable only for modular switch systems.

4.2.9.2.2 chassis ha

	<code>chassis ha reset other</code> Performs a reset to the other management card in the chassis.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Example	<code>switch (config) # chassis ha reset other</code>
Related Commands	show chassis ha
Notes	This command is applicable only for modular switch systems.

4.2.9.2.3 chassis ha power enable other

	<code>chassis ha power enable other</code> <code>no chassis ha power enable other</code> Enables the other management card in the chassis. The no form of the command disables the other management card in the chassis.
Syntax Description	N/A
Default	The other management card is enabled
Configuration Mode	config
History	3.1.0000
Example	<code>switch (config) # chassis ha power enable other</code>
Related Commands	show chassis ha
Notes	This command is applicable only for modular switch systems.

4.2.9.2.4 show chassis ha

	<code>show chassis ha</code> Displays chassis HA parameters and status.
Syntax Description	N/A
Default	The other management card is enabled
Configuration Mode	Any command mode
History	3.1.0000

Example	<pre>switch (config) # show chassis ha 2-node HA state: Box management IPv4: 10.7.146.44/24 Box management IPv6: fdfd:fdfd:7:145::1033:47fd/64 interface : mgmt0 local role : master local slot : 1 other state : ready reset count : 0</pre>
Related Commands	chassis ha
Notes	This command is applicable only for modular switch systems.

4.2.9.2.5 chassis ha bipv6

	chassis ha bipv6 {ipv6 address} {ipv6 mask length} [force] The command configures the Box IPv6.	
Syntax Description	ipv6 address	The ipv6 box ip
	ipv6 mask length	The mask for IPv6 box ip
Default	The other management card is enabled	
Configuration Mode	Any command mode	
History	3.8.1200	
Example	<pre>switch (config) # chassis ha bipv6 fdfd:fdfd:7:145::1033:47fd /64</pre>	
Related Commands	chassis ha	
Notes		

4.3 UNBREAKABLE-LINK® Adapter and Switch Technology



NVIDIA adapter, switch, and interconnect products support a unique UNBREAKABLE-LINK® technology that ensures the network links stay connected (operational) even in a bad connectivity environment.

MLNX-OS offers PHY profile configuration for InfiniBand interfaces. PHY profile includes Link Level Retransmission (LLR) configuration. A PHY profile is bound to any InfiniBand interface.

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

LLR transmitter breaks the transmitted Layer 2 data stream into Cells and adds a CRC checksum to each cell.

LLR receiver checks the Cell CRC, in case there is no CRC errors, it forwards the cell and acknowledges the peer.

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

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

LLR is not operational for cables longer than 30m.

4.3.1 LLR Mode

The following LLR modes are applicable per port per speed:

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

4.3.2 LLR Negotiation

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

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

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

4.3.3 LLR Status

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

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

4.3.4 UNBREAKABLE-LINK® Switch Commands

4.3.4.1 show interfaces ib llr

	show interfaces ib [<number>] llr Displays LLR status	
Syntax Description	number	The interface number
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.0500	

Example	<pre>switch (config) # show interfaces ib llr ----- Interface LLR status ----- IB1/1 Inactive IB1/2 Inactive IB1/3 Inactive IB1/4 Inactive IB1/5 Inactive IB1/6 Inactive IB1/7 Inactive IB1/8 Inactive IB1/9 Inactive IB1/10 Inactive IB1/11 Inactive IB1/12 Inactive IB1/13 Inactive ... </pre>
Related Commands	
Notes	

4.4 Upgrade/Downgrade Process

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

- [Important Pre-OS Upgrade Notes](#)
- [Upgrading Operating System Software](#)
- [Upgrading HA Groups](#)
- [Upgrading MLNX-OS Software on Modular Switches](#)
- [Deleting Unused Images](#)
- [Downgrading OS Software](#)
- [Upgrading System Firmware](#)
- [Software Management Commands](#)

4.4.1 Important Pre-OS Upgrade Notes

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

- Upgrading modular switch systems can take up to 30 minutes during which time the system is indisposed
- Upgrading the OS while embedded SM is enabled may cause the command “no hostname” to fail upon first execution. To resolve this, rerun the command
- The upgrade procedure burns the software image as well as the firmware should there be a need
- Before upgrading the software image on your system, make sure to close all CLI sessions besides the one used to run the upgrade process
- If running a system with dual management cards, refer to [“Upgrading MLNX-OS Software on Modular Switches”](#)
- To upgrade the MLNX-OS version on an SM cluster, please refer to [“Upgrading HA Groups”](#)
- The End-User License Agreement (EULA) must read and accepted after image upgrade in case the EULA is modified. The EULA link is only available upon first login to CLI

- Linux docker container names are limited to 180 characters. Upgrading to this version removes containers which do not comply with this limitation and prints the following warning to the log: “Removed configuration of container: <container name>, container name is limited to 180 characters”

4.4.2 Upgrading Operating System Software

To upgrade MLNX-OS, perform the following steps.

1. Enter Config mode.

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

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

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

  Partition 1:
  <old_image>

  Partition 2:
  <old_image>

Last boot partition: 1
Next boot partition: 1

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

Serve image files via HTTP/HTTPS: no
No image install currently in progress.

Boot manager password is set.

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

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

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

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

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

4. Fetch the new software image.

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

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

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

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

  Partition 1:
  <old_image>

  Partition 2:
  <old_image>

Last boot partition: 1
Next boot partition: 1

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

Serve image files via HTTP/HTTPS: no

No image install currently in progress.

Boot manager password is set.

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

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

6. Install the new image.

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

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

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

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

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

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

  Partition 1:
  <new_image>

  Partition 2:
  <old_image>

Last boot partition: 1
Next boot partition: 1

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

Serve image files via HTTP/HTTPS: no

No image install currently in progress.

Boot manager password is set.

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

Settings for next boot only:
```



```
Fallback reboot on configuration failure: yes (default)
```

9. Save current configuration.

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

10. Reboot to run the new image.

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

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

On systems with dual management, the software must be upgraded on both the host and the device modules.

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

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

4.4.3 Upgrading HA Groups

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

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

1. Identify the HA group master.

For IB HA. Run:

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

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

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

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

4.4.4 Upgrading MLNX-OS Software on Modular Switches

Modular switches feature dual management modules.

1. Identify the chassis HA master. Run:

```
show chassis ha
```

2. Upgrade the chassis master according to steps 1-8 in "[Upgrading Operating System Software](#)". Please DO NOT reboot!
3. Upgrade the second management module according to steps 1-8 in "[Upgrading Operating System Software](#)". Please DO NOT reboot!
4. Reset the slave management module. In the master management module, run:

```
chassis ha reset other
```

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

```
reload force immediate
```

An alternative for steps 4 and 5 is to power cycle the system.

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

```
show chassis ha
```

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

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

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

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

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

4.4.5 Deleting Unused Images

To delete unused images, conduct the following steps.

1. Get a list of the unused images.

```
switch (config) # show images
Installed images:
  Partition 1:
    version: image-X86_64-3.6.5000.img
  Partition 2:
    version: image-X86_64-3.6.5000.img
Last boot partition: 1
Next boot partition: 1
Images available to be installed:
  No image files are available to be installed.
Serve image files via HTTP/HTTPS: no
No image install currently in progress.
Boot manager password is set.
Image signing           : trusted signature always required
Admin require signed images: yes
Settings for next boot only:
  Fallback reboot on configuration failure: yes (default)
```

2. Delete the unused images.

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

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

4.4.6 Downgrading OS Software



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

1. Log in to the switch via the CLI using the console port.
2. Backup configuration by following these steps.
 - a. Disable paging of CLI output.

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

- b. Display commands to recreate current running configuration.

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

- c. Copy the output to a text file.

4.4.6.1 Downloading Image

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

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

2. Log in to [NVIDIA Enterprise Support Portal](#) and download the relevant MLNX-OS version to your system type
3. Log in to your system via the CLI.
4. Change to Config mode.

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

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

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

4.4.6.2 Downgrading Image

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

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

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

3. Display all image files on the system.

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

4. Install the fetched image.

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

5. Display all image files on the system.

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

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

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

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

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

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

7. Reload.

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

4.4.6.3 Switching to Partition with Older Software Version

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

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

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

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

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

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

1. Run the following command.

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

2. Set the boot partition.

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

3. Save the configuration.

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

4. Reload the system.

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

4.4.7 Upgrading System Firmware

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

4.4.7.1 After Updating Software

Upon rebooting your switch system after updating the OS software, the OS compares its default firmware version with the currently programmed firmware versions on all the switch modules (leafs and spines on modular-class switches, or simply the switch card on modular switch systems). If one or more of the switch modules is programmed with a firmware version other than the default version, then the OS automatically attempts to burn the default firmware version instead.

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

To verify that the firmware update was successful, log into your switch and run the command “show asic-version” (can be run in any mode). This command lists all of the switch modules along with their firmware versions. Make sure that all the firmware versions are the same and match the default firmware version. If the firmware update failed for one or more modules, then the following warning is displayed.

Some subsystems are not updated with a default firmware.

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

4.4.7.2 After Inserting a Switch Spine or Leaf

This section is applicable to modular switch systems only.

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

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

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

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

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

4.4.7.3 Importing Firmware and Changing the Default Firmware

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

4.4.7.3.1 Default Firmware Change on Standalone Systems

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

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

2. Save the configuration.

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

3. Reboot the system to enable auto update.

4.4.7.3.2 Default Firmware Change Dual Management Systems

This flow should be implemented on both management modules in parallel.

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

```
switch (config) # image fetch scp://username:password@10.7.34.115//my_directory/fw-SIB-rel-11_1600_0200-
FIT.mfa
100.0% [#####]
```

2. Change default firmware on the management modules using the command `image default-chip-fw`.
3. Verify that both master and slave have successfully installed the new firmware. The following message should be displayed:

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

4. Run "configuration write" on both management modules.
5. Run "chassis ha reset other" on the master management module only.
6. Run "reload" on the master management module.

4.4.8 Software Management Commands



- [4.4.8.1 image boot](#)
- [4.4.8.2 boot next](#)
- [4.4.8.3 boot system](#)
- [4.4.8.4 image default-chip-fw](#)
- [4.4.8.5 image delete](#)
- [4.4.8.6 image fetch](#)
- [4.4.8.7 image install](#)
- [4.4.8.8 image move](#)
- [4.4.8.9 image options](#)
- [4.4.8.10 show bootvar](#)
- [4.4.8.11 show images](#)

4.4.8.1 image boot

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

4.4.8.2 boot next

	boot next fallback-reboot enable no boot next fallback-reboot enable Sets the default setting for next boot. Normally, if the system fails to apply the configuration on startup (after attempting upgrades or downgrades, as appropriate), it will reboot to the other partition as a fallback. The no form of the command tells the system not to do that, only for the next boot.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config	
History	3.2.0506	
Example	switch (config) # boot next fallback-reboot enable	
Related Commands	show images	
Notes	<ul style="list-style-type: none"> • Normally, if the system fails to apply the configuration on startup (after attempting upgrades or downgrades, as appropriate) it reboots to the other partition as a fallback. • The no form of this command tells the system not to do that only for the next boot. In other words, this setting is not persistent and goes back to being enabled automatically after each boot. • When downgrading to an older software version which has never been run yet on a system, the “fallback reboot” always happens, unless the command “no boot next fallback-reboot enable” is used. However, this also happens when the older software version has been run before, but the configuration file has been switched since upgrading. In general, a downgrade only works (without having the fallback reboot forcibly disabled) if the process can find a snapshot of the configuration file (by the same name as the currently active one) which was taken before upgrading from the older software version. If that is not found, a fallback reboot is performed in preference to falling back to the initial database because the latter generally involves a loss of network connectivity, and avoiding that is of paramount importance. 	

4.4.8.3 boot system

	boot system {location next} no boot system next Configures which system image to boot by default. The no form of the command resets the next boot location to the current active one.	
Syntax Description	location	Specifies location from which to boot system <ul style="list-style-type: none"> • 1—installs to location 1 • 2—installs to location 2
	next	Boots system from next location after one currently booted
Default	N/A	
Configuration Mode	config	
History	3.2.0506	
Example	switch (config) # boot system location 2	
Related Commands	show images	

Notes	
-------	--

4.4.8.4 image default-chip-fw

	<pre>image default-chip-fw <filename></pre> <pre>no image default-chip-fw <original-fw-filename></pre> <p>Sets the default firmware package to be installed. The no form of the command resets default firmware package.</p>	
Syntax Description	filename	Specifies the firmware filename
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.6.6000	Added the no form of the command
Example	switch (config) # image default-chip-fw <filename>.mfa	
Related Commands	show asic-version show images	
Notes		

4.4.8.5 image delete

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

4.4.8.6 image fetch

	<pre>image fetch <URL> [<filename>]</pre> <p>Downloads an image from the specified URL or via SCP.</p>	
Syntax Description	URL	HTTP, HTTPS, FTP, TFTP, SCP and SFTP are supported Example: scp://username[:password]@hostname/path/filename
	filename	Specifies a filename for this image to be stored as locally
Default	N/A	
Configuration Mode	config	

History	3.1.0000 3.9.2000—Added VRF option
Example	<pre>switch (config) # image fetch scp://<username>@192.168.10.125/var/www/html/<image_name> Password ***** 100.0%[#####] switch (config) # Other options: switch (config) # image fetch http://10.1.0.40/path/filename switch (config) # image fetch http://[fd4f:13:cc00:1::40]/path/filename switch (config) # image fetch ftp://user:mypassword@10.1.0.40/foo/bar.img switch (config) # image fetch ftp://user:mypassword@[fd4f:13:cc00:1::40]/foo/bar.img switch (config) # image fetch tftp://hostname/dir/filename switch (config) # image fetch tftp://[fd4f:13:cc00:1::40]/dir/filename switch (config) # image fetch scp://user@myhost/dir/filename switch (config) # image fetch scp://user@myhost:1022/dir/filename switch (config) # image fetch scp://user:pass@[fd4f:13:cc00:1::40]/dir/filename switch (config) # image fetch sftp://user@myhost/dir/filename switch (config) # image fetch sftp://user@[fd4f:13:cc00:1::40]:1022/dir/filename switch (config) # image fetch sftp://user:pass@[fd4f:13:cc00:1::40]/dir/filename</pre>
Related Commands	show images
Notes	<ul style="list-style-type: none"> • Please delete the previously available image, prior to fetching the new image • The path to the file in the case of TFTP depends on the server configuration. Therefore, it may not be an absolute path but a relative one. • See "Upgrading Operating System Software" page

4.4.8.7 image install

	image install <image-filename> [location <location-ID>] [progress <prog-options>] Installs the specified image file.	
Syntax Description	image-filename	Specifies the image name
	location-ID	Specifies the image destination location
	prog-options	<ul style="list-style-type: none"> • “no-track” overrides CLI default and does not track the installation progress • “track” overrides CLI default and tracks the installation progress
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # image install X86_64 3.6.5000 2017-07-26 06:54:12 x86_64 Step 1 of 4: Verify Image 100.0% [#####] Step 2 of 4: Uncompress Image 100.0% [#####] Step 3 of 4: Create Filesystems 100.0% [#####] Step 4 of 4: Extract Image 100.0% [#####] switch (config) #</pre>	
Related Commands	show images	

Notes	<ul style="list-style-type: none"> • The image cannot be installed on the “active” location (the one which is currently being booted) • On a two-location system, the location is chosen automatically if no location is specified
-------	--

4.4.8.8 image move

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

4.4.8.9 image options

	<code>image options serve all</code> <code>no image options serve all</code> Configures options and defaults for image usage. The no form of the command disables options and defaults for image usage.	
Syntax Description	<code>serve all</code>	Specifies that the image files present on this appliance should be made available for HTTP and/or HTTPS download
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	<code>switch (config) # image options serve all</code>	
Related Commands	show images	
Notes	The parameter “serve all” affects not only the files currently present, but also any files that are later downloaded. It only applies to image files, not the installed images, which are not themselves in a downloadable format. After running “serve all” the URLs where the images will be available are: <ul style="list-style-type: none"> • <code>http://<HOSTNAME>/system_images/<FILENAME></code> • <code>https://<HOSTNAME>/system_images/<FILENAME></code> 	

4.4.8.10 show bootvar

	<code>show bootvar</code> Displays the installed system images and the boot parameters.
Syntax Description	N/A

Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config)# show bootvar Installed images: Partition 1: X86_64 3.6.4110-12 2017-07-26 06:54:12 x86_64 Partition 2: X86_64 3.6.4006 2017-07-03 16:17:39 x86_64 Last boot partition: 1 Next boot partition: 1 Serve image files via HTTP/HTTPS: no Boot manager password is set. Image signing: trusted signature always required Admin require signed images: yes Settings for next boot only: Fallback reboot on configuration failure: yes (default)</pre>
Related Commands	
Notes	

4.4.8.11 show images

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

4.5 Configuration Management



4.5.1 Saving a Configuration File

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

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

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

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

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

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

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

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

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

4.5.2 Loading a Configuration File

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

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

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

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

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

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

4.5.3 Restoring Factory Default Configuration

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

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

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

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

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

- a. Connect to a remote console/serial connection.
- b. Remove the slave management module.
- c. Run "reset factory":

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

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

- d. Log in as "admin" and start running the Configuration Wizard.
- e. Insert the slave management module.
- f. Remove the master management module.

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

- g. Repeat Step 3 on the new Master management module.
- h. Insert the other management module. No takeover will occur at this stage.
- i. Power cycle the system.

4.5.4 Managing Configuration Files

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

4.5.4.1 BIN Configuration Files

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

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

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

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

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

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

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

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

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

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

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

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

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

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

4.5.4.2 Text Configuration Files

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

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

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

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

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

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

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


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

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

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

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

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

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

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

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

4.5.5 Automated Periodic Configuration File Backup

4.5.5.1 Automated Backup

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

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

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

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

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

- To save the configuration, run the following:

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

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

- To remove the remote URL, run the following:

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

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

4.5.5.2 Automated Periodic Backup

Scheduled jobs can be used to perform automated periodic backup.

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

1. Create a job.

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

2. Add the upload command to the job.

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

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

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

4. Enable the job.

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

4.5.6 Configuration Management Commands



- [4.5.6.1 File System](#)
 - [4.5.6.1.1 debug generate dump](#)
 - [4.5.6.1.2 file debug-dump](#)
 - [4.5.6.1.3 file stats](#)
 - [4.5.6.1.4 file tcpdump](#)
 - [4.5.6.1.5 file eula upload](#)
 - [4.5.6.1.6 file open-source-licenses upload](#)
 - [4.5.6.1.7 file help-docs upload](#)
 - [4.5.6.1.8 reload](#)
 - [4.5.6.1.9 reset factory](#)
 - [4.5.6.1.10 configuration new factory](#)
 - [4.5.6.1.11 configuration new factory keep-docker](#)
 - [4.5.6.1.12 show files debug-dump](#)
 - [4.5.6.1.13 show files stats](#)
 - [4.5.6.1.14 show files system](#)
 - [4.5.6.1.15 show files tcpdump](#)
- [4.5.6.2 Configuration Files](#)
 - [4.5.6.2.1 configuration audit](#)

- [4.5.6.2.2 configuration auto-upload](#)
- [4.5.6.2.3 configuration copy](#)
- [4.5.6.2.4 configuration delete](#)
- [4.5.6.2.5 configuration fetch](#)
- [4.5.6.2.6 configuration jump-start](#)
- [4.5.6.2.7 configuration merge](#)
- [4.5.6.2.8 configuration move](#)
- [4.5.6.2.9 configuration new](#)
- [4.5.6.2.10 configuration revert](#)
- [4.5.6.2.11 configuration switch-to](#)
- [4.5.6.2.12 configuration text fetch](#)
- [4.5.6.2.13 configuration text file](#)
- [4.5.6.2.14 configuration text generate](#)
- [4.5.6.2.15 configuration upload](#)
- [4.5.6.2.16 configuration write](#)
- [4.5.6.2.17 write](#)
- [4.5.6.2.18 show configuration](#)
- [4.5.6.2.19 show configuration auto-upload](#)
- [4.5.6.2.20 show running-config](#)
- [4.5.6.2.21 show running-config interface](#)

4.5.6.1 File System

4.5.6.1.1 debug generate dump

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

4.5.6.1.2 file debug-dump

	file debug-dump {delete {<filename> all latest} email {<filename> latest} upload {<filename> latest} <URL>}
	Manipulates debug dump files.

Syntax Description	delete	Deletes a debug dump file. <ul style="list-style-type: none"> all—deletes all existing debug files from this machine latest—deletes latest debug file from this machine
	email	Emails a debug dump file to pre-configured recipients for “informational events”. <ul style="list-style-type: none"> latest—emails the latest debug file to a pre-configured recipients
	upload	Uploads a debug dump file to a remote host. <ul style="list-style-type: none"> latest—uploads the latest debug file to a remote host
	URL	The URL to the remote host. Supported URL formats: HTTP, HTTPS, FTP, TFTP, SCP and SFTP. Example: scp://username[:password]@hostname/path/filename
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.3.4000	Added “all” and “latest” options
Example	switch (config) # file debug-dump email sysdump-switch-112104-20114052-091707.tgz	
Related Commands	show files debug-dump	

4.5.6.1.3 file stats

	file stats {delete <filename> move {<source filename> <destination filename>} upload <filename> <URL>} Manipulates statistics report files.	
Syntax Description	delete <filename>	Deletes a stats report file.
	move <source filename> <destination filename>	Renames a stats report file.
	upload <filename> <URL>	Uploads a stats report file. Supported URL formats: HTTP, HTTPS, FTP, TFTP, SCP and SFTP. Example: scp://username[:password]@hostname/path/filename
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # file stats move memory-1.csv memory-2.csv	
Related Commands	show files stats show files stats <filename>	
Notes		

4.5.6.1.4 file tcpdump

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

4.5.6.1.5 file eula upload

	file eula upload <filename> <URL> Uploads the End User License Agreement to a specified remote location.	
Syntax Description	filename	The End User License Agreement
	URL	URL or scp://username[:password]@hostname/path/filename
Default	N/A	
Configuration Mode	config	
History	3.4.1100	
Example	switch (config) # file eula upload MLNX-OS_EULA.pdf ? <URL or scp://username[:password]@hostname/path/filename >	
Related Commands	license	
Notes	N/A	

4.5.6.1.6 file open-source-licenses upload

	file open-source-licenses upload <filename> <URL> Uploads the Open Source Licenses file.	
Syntax Description	filename	The Open Source Licenses file
	URL	URL or scp://username[:password]@hostname/path/filename
Default	N/A	
Configuration Mode	config	

History	3.9.3100
Example	<code>switch (config) # file open-source-licenses upload Open_Source_Licenses.txt scp://username[:password]@hostname/path/filename</code>
Related Commands	license
Notes	N/A

4.5.6.1.7 file help-docs upload

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

4.5.6.1.8 reload

	<code>reload [force immediate halt [noconfirm] noconfirm]</code> 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	
Example	<code>switch (config) # reload</code> Configuration has been modified; save first? [yes] yes Configuration changes saved. ...	
Related Commands	reset factory	
Notes		

4.5.6.1.9 reset factory

	reset factory [keep-all-config keep-basic keep-config-group keep-virt-vols keep-docker keep-docker clear-label <label name>] only-config] [halt] Clears the system and resets it entirely to its factory state.	
Syntax Description	keep-all-config	Preserves all configuration files including licenses. Removes the logs, stats, images, snapshots, history, and known hosts. The user is prompted for confirmation before honoring this command, unless confirmation is disabled with the command: “no cli default prompt confirm-reset”.
	keep-basic	Preserves licenses in the running configuration file.
	keep-config-group	Reset to the factory defaults of the current RoCE config group: no-roce, lossless, lossy or semi-lossless.
	keep-virt-vols	Preserves all virtual disk volumes.
	only-config	Removes configuration files only. Logs, stats, images, snapshots, history, and known hosts are preserved.
	halt	The system is halted after this process completes.
	keep-docker	Preserves all current docker configurations.
	keep-docker clear-label <label name>	Preserves all current docker configurations, but deletes the content of the given docker storage label. (Note that only the content of the label folder will be deleted. The label itself will remain intact.)
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.4.0000	Added notes and “keep-virt-vols” parameter
	3.6.2002	Updated example and notes
	3.8.1300	Added “keep-docker” and “keep-docker clear-label” option
Example	<pre>switch (config) # reset factory Warning - confirming will cause system reboot. Type 'YES' to confirm reset: YES Resetting and rebooting the system -- please wait... ...</pre>	
Related Commands	reload	
Notes	<ul style="list-style-type: none"> • Effects of parameter “keep-all-config”: Licenses—not deleted; profile—no change; configuration—unchanged; management IP—unchanged • Effects of parameter “keep-basic”: Licenses—not deleted; profile—reset; configuration—reset; management IP—reset • Effects of parameter “keep-virt-vols”: Licenses—deleted; profile—reset; configuration—reset; management IP—deleted • Confirming the command causes system reboot 	

4.5.6.1.10 configuration new factory

	configuration new <filename> factory Creates new file with only factory defaults.
Syntax Description	N/A

Default	N/A
Configuration Mode	config
History	3.7.1102
Example	switch (config) # no configuration new my_file factory
Related Commands	configuration new factory configuration new factory keep-basic configuration new factory keep-connect
Notes	

4.5.6.1.11 configuration new factory keep-docker

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

4.5.6.1.12 show files debug-dump

	show files debug-dump [<filename>] Displays a list of debug dump files.
Syntax Description	filename Displays a summary of the contents of a particular debug dump file.
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config) # show files debug-dump sysdump-switch-20170731-161038.tgz ===== System information: Hostname: switch Version: X86_64 3.6.4006 2017-07-03 16:17:39 x86_64 Current time: 2017-07-31 16:10:38 System uptime: 19d 18h 20m 12s ===== Output of 'uname -a': Linux switch 3.10.0-327.36.3.el7smp-x86_64 X86_64 jenkins #1 2017-06-27 12:34:55 SMP x86_64 x86_64 x86_64 GNU/Linux =====
Related Commands	file debug-dump

Notes	
-------	--

4.5.6.1.13 show files stats

	show files stats <filename> Displays a list of statistics report files.	
Syntax Description	filename	Display the contents of a particular statistics report file.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example	<pre>switch (config) # show files stats memory-201140524-111745.csv</pre>	
Related Commands	file stats	
Notes		

4.5.6.1.14 show files system

	show files system [detail] Displays usage information of the file systems on the system.	
Syntax Description	detail	Displays more detailed information on file-system.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example	<pre>switch (config) # show files stats memory-201140524-111745.csv</pre>	
Related Commands		
Notes		

4.5.6.1.15 show files tcpdump

	show files tcpdump Displays a list of statistics report files.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example	<pre>switch (config) # show files stats test dump3</pre>	
Related Commands		
Notes		

4.5.6.2 Configuration Files

4.5.6.2.1 configuration audit

	configuration audit max-changes <number> Chooses settings related to configuration change auditing.	
Syntax Description	max-changes	Set maximum number of audit messages to log per change.
Default	1000	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # configuration audit max-changes 100	
Related Commands	show configuration	
Notes		

4.5.6.2.2 configuration auto-upload

	configuration auto-upload remote-url no configuration auto-upload remote-url Sets the remote URL to upload for automated backup. The no form resets the remote URL.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config	
History	3.9.0500	
Example	switch (config) # configuration auto-upload remote-url "scp:// root:password@192.168.10.125/tmp/conf1"	
Related Commands	show configuration auto-upload	
Notes	If this feature is set, after every configuration write it will upload the active configuration file to the configured remote URL.	

4.5.6.2.3 configuration copy

	configuration copy <source-name> <dest-name> Copies a configuration file.	
Syntax Description	source-name	Name of source file.
	dest-name	Name of destination file. If the file of specified filename does not exist a new file will be created with said filename.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # configuration copy initial.bak example	

Related Commands	
Notes	<ul style="list-style-type: none"> • This command does not affect the current running configuration • The active configuration file may not be the target of a copy. However, it may be the source of a copy in which case the original remains active.

4.5.6.2.4 configuration delete

	configuration delete <filename> Deletes a configuration file.	
Syntax Description	filename	Name of file to delete
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # configuration delete example	
Related Commands	show configuration files	
Notes	<ul style="list-style-type: none"> • This command does not affect the current running configuration • The active configuration file may not be deleted 	

4.5.6.2.5 configuration fetch

	configuration fetch <URL> [<name>] Downloads a configuration file from a remote host.	
Syntax Description	URL	Supported formats: HTTP, HTTPS, FTP, TFTP, SCP and SFTP. Example: scp://username[:password]@hostname/path/filename
	name	The name of the configuration file.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # configuration fetch scp://root:password@192.168.10.125/tmp/ conf1	
Related Commands	configuration switch-to	
Notes	<ul style="list-style-type: none"> • The downloaded file should not override the active configuration file, using the <name> parameter • If no name is specified for a configuration fetch, it is given the same name as it had on the server • No configuration file may have the name “active” 	

4.5.6.2.6 configuration jump-start

	configuration jump-start Runs the initial-configuration wizard.	
Syntax Description	N/A	
Default	N/A	

Configuration Mode	config
History	3.1.0000
Example	<pre>switch (config) # configuration jump-start Configuration wizard Step 1: Hostname? [switch-3cc29c] Step 2: Use DHCP on mgmt0 interface? y Step 3: Admin password (Enter to leave unchanged)? You have entered the following information: 1. Hostname: switch-3cc29c 2. Use DHCP on mgmt0 interface: yes 3. Enable IPv6: yes 4. Enable IPv6 autoconfig (SLAAC) on mgmt0 interface: yes 53. Admin password (Enter to leave unchanged): (unchanged) To change an answer, enter the step number to return to. Otherwise hit <enter> to save changes and exit. Choice: Configuration changes saved.</pre>
Related Commands	configuration switch-to
Notes	<ul style="list-style-type: none"> 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)

4.5.6.2.7 configuration merge

	configuration merge <filename> Merges the “shared configuration” from one configuration file into the running configuration.	
Syntax Description	filename	Name of file from which to merge settings.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # configuration merge new-config-file</pre>	
Related Commands		
Notes	<ul style="list-style-type: none"> No configuration files are modified during this process The configuration filename must be a non-active configuration file 	

4.5.6.2.8 configuration move

	configuration move <source-name> <dest-name> Renames a configuration file.	
Syntax Description	source-name	Name of file to rename.
	dest-name	New name of renamed file.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	<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</pre>	

Related Commands	show configuration
Notes	<ul style="list-style-type: none"> • This command does not affect the current running configuration • The active configuration file may not be the target of a move

4.5.6.2.9 configuration new

	<code>configuration new <filename> [factory [keep-basic] [keep-connect]]</code> Creates a new configuration file under the specified name. The parameters specify what configuration, if any, to carry forward from the current running configuration.	
Syntax Description	filename	Names for new configuration file.
	factory	Creates new file with only factory defaults.
	keep-basic	Keeps licenses and host keys.
	keep-connect	Keeps configuration necessary for connectivity (interfaces, routes, and ARP).
Default	Keeps licenses and host keys	
Configuration Mode	config	
History	3.1.0000	
Example	<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</pre>	
Related Commands	show configuration	
Notes	<ul style="list-style-type: none"> • This command does not affect the current running configuration • The active configuration file may not be the target of a move 	

4.5.6.2.10 configuration revert

	<code>configuration revert {factory [keep-basic keep-connect] saved}</code> Reverts the system configuration to a previous state.	
Syntax Description	factory	Creates new file with only factory defaults.
	keep-basic	Keeps licenses and host keys.
	keep-connect	Keeps configuration necessary for connectivity (interfaces, routes, and ARP).
	saved	Reverts running configuration to last saved configuration.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # configuration revert saved</pre>	
Related Commands	show configuration	
Notes	<ul style="list-style-type: none"> • This command is not available on IB multi-SWID system profile 	

4.5.6.2.11 configuration switch-to

	configuration switch-to <filename>[no-reboot] Loads the configuration from the specified file and makes it the active configuration file.
Syntax Description	no-reboot Forces configuration change without rebooting.
Default	N/A
Configuration Mode	config
History	3.1.0000 3.6.1002 Added “no-reboot” option
Example	<pre>switch (config) # show configuration files initial (active) newcon initial.prev initial.bak switch (config) # configuration switch-to newcon no-reboot switch (config) # show configuration files initial newcon (active) initial.prev initial.bak</pre>
Related Commands	show configuration files
Notes	<ul style="list-style-type: none"> • The current running configuration is lost and not automatically saved to the previous active configuration file • When running the command without the “no-reboot” parameter, the user is prompted to OK a reboot. If the answer is “yes”, the configuration is replaced and the system is rebooted immediately

4.5.6.2.12 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	<table border="1"> <tr> <td>apply</td> <td>Applies the file to the running configuration (i.e. executes the commands in it). This option has the following parameters: <ul style="list-style-type: none"> • 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 </td> </tr> <tr> <td>filename</td> <td>Specifies filename for saving downloaded text file.</td> </tr> <tr> <td>overwrite</td> <td>Downloads the file and saves it using the same name it had on the server. This option has the following parameters: <ul style="list-style-type: none"> • apply—applies the downloaded configuration to the running system • filename—specifies filename for saving downloaded text file </td> </tr> </table>	apply	Applies the file to the running configuration (i.e. executes the commands in it). This option has the following parameters: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • apply—applies the downloaded configuration to the running system • filename—specifies filename for saving downloaded text file
	apply	Applies the file to the running configuration (i.e. executes the commands in it). This option has the following parameters: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • apply—applies the downloaded configuration to the running system • filename—specifies filename for saving downloaded text file 						
Default	N/A						
Configuration Mode	config						

History	3.2.1000
Example	<code>switch (config) # configuration text fetch scp://username[:password]@hostname/path/filename</code>
Related Commands	
Notes	

4.5.6.2.13 configuration text file

	<code>configuration text file <filename> {apply [fail-continue] [verbose] [reboot] delete rename <filename> upload <URL>}</code> Performs operations on text-based configuration files.	
Syntax Description	<code>filename <file></code>	Specifies the filename.
	<code>apply</code>	Applies the configuration on the system.
	<code>fail-continue</code>	Continues execution of the commands even if some commands fail.
	<code>verbose</code>	Displays all commands being executed and their output, instead of just those that get errors.
	<code>delete</code>	Deletes the file.
	<code>rename <filename></code>	Renames the file.
	<code>upload <URL></code>	Supported types are HTTP, HTTPS, FTP, TFTP, SCP and SFTP. For example: <code>scp://username[:password]@hostname/path/filename</code>
	<code>reboot</code>	Write the configuration and reboot after successful execution.
Default	N/A	
Configuration Mode	config	
History	3.1.0000 3.9.0300 Added ability to apply reboot	
Example	<code>switch (config) # configuration text file my-config-file delete</code>	
Related Commands	show configuration files	
Notes		

4.5.6.2.14 configuration text generate

	<code>configuration text generate {active {running saved} file <filename> } {save <filename> upload <URL>}</code> Generates a new text-based configuration file from this system's configuration.	
Syntax Description	<code>active</code>	Generates from currently active configuration.
	<code>running</code>	Uses running configuration.
	<code>saved</code>	Uses saved configuration.
	<code>file <filename></code>	Generates from inactive saved configuration.
	<code>save</code>	Saves new file to local persistent storage.

	upload <URL>	Supported types are HTTP, HTTPS, FTP, TFTP, SCP and SFTP. For example: scp://username[:password]@hostname/path/filename.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # configuration text generate file initial.prev save example	
Related Commands	show configuration files	
Notes		

4.5.6.2.15 configuration upload

	configuration upload {timestamp} {active <name>} <URL or scp or sftp://username:password@hostname[:port]/path/filename> Uploads a configuration file to a remote host.	
Syntax Description	active	Upload the active configuration file.
	timestamp	Will append the timestamp to the filename uploaded to remote.
Default	N/A	
Configuration Mode	config	
History	3.1.0000 3.9.0500 Added timestamp option	
Example	switch (config) # configuration upload active scp://root:password@192.168.10.125/tmp/conf1	
Related Commands	show configuration files	
Notes	No configuration file may have the name “active” or “timestamp”.	

4.5.6.2.16 configuration write

	configuration write [local to <filename> [no-switch]] Saves the running configuration to the active configuration file.	
Syntax Description	local	Saves the running configuration locally (same as “write memory local”).
	to <filename>	Saves the running configuration to a new file under a different name and makes it the active file.
	no-switch	Saves the running configuration to this file but keep the current one active.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # configuration write	
Related Commands	write	

Notes	
-------	--

4.5.6.2.17 write

	write {memory [local] terminal} Saves or displays the running configuration.	
Syntax Description	memory	Saves running configuration to the active configuration file. It is the same as “configuration write”.
	local	Saves the running configuration only on the local node. It is the same as “configuration write local”.
	terminal	Displays commands to recreate current running configuration. It is the same as “show running-config”.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	<pre> switch (config) # write terminal ## ## Running database "initial" ## Generated at 2014/05/27 10:05:16 +0000 ## Hostname: switch ## ## ## Network interface configuration ## interface mgmt0 comment "" interface mgmt0 create interface mgmt0 dhcp interface mgmt0 display interface mgmt0 duplex auto interface mgmt0 mtu 1500 no interface mgmt0 shutdown interface mgmt0 speed auto no interface mgmt0 zeroconf ## ## Local user account configuration ## username a** capability admin no username a** disable username a** disable password </pre>	
Related Commands	show running-config configuration write	
Notes		

4.5.6.2.18 show configuration

	show configuration [audit files [<filename>] running text files] Displays a list of CLI commands that will bring the state of a fresh system up to match the current persistent state of this system.	
Syntax Description	audit	Displays settings for configuration change auditing.

	files [<filename>]	Displays a list of configuration files in persistent storage if no filename is specified. If a filename is specified, it displays the commands to recreate the configuration in that file. In the latter case, only non-default commands are shown, as for the normal “show configuration” command.
	running	Displays commands to recreate current running configuration. Same as the command “show configuration” except that it applies to the currently running configuration, rather than the current persisted configuration.
	text files	Displays names of available text-based configuration files.
Default	N/A	
Configuration Mode	config	
History	3.1.0000 3.3.5006 Removed “running full” and “full” parameters	
Example	<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</pre>	
Related Commands		
Notes		

4.5.6.2.19 show configuration auto-upload

	show configuration auto-upload Shows the automated backup settings.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.9.0500
Example	<pre>switch (config) # show configuration auto-upload Auto-upload settings: Enabled: yes Remote url: scp://root@192.168.10.125/tmp/conf1 Password : *****</pre>
Related Commands	configuration auto-upload remote-url
Notes	If this feature is set. After every configuration write, it will upload the active configuration file to the configured remote URL.

4.5.6.2.20 show running-config

	show running-config [expanded protocol <protocol> diff diff <config_file_name>] Displays commands to recreate current running configuration.	
Syntax Description	expanded	Displays commands in expanded format without compressing ranges.
	protocol	Only displays commands relating to the specified protocol.
	diff	Displays delta between saved config file (active by default) and running-config.
	config_file_name	Displays delta between the specified saved config file and running-config.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.3.4402	Removed “full” parameter
	3.6.2002	Updated example and added parameters
	3.6.3640	Added support for forwarding mode configuration
	3.8.1000	Added support to show diff between running-config and saved config files (active file saved by default)
Example		
<pre>switch (config) # show running-config diff Only in running-config: + interface port-channel 1 + interface ethernet 1/31-1/33 speed 10G force + interface port-channel 1 description lag Only in saved configuration file: - ip route 169.254.22.0/24 169.254.2.100 Common configuration but in different order in saved configuration file and running-config: <<None>></pre>		
Related Commands		
Notes	<ul style="list-style-type: none"> + <string> : <string> exists only in running-config, but not in the saved filename (or active config file if no <filename> is specified) - <string> : <string> does not exist in running-config, but exists in the saved filename (or active config file if no <filename> is specified) ! <string> : <string> exists in both running-config and the saved filename, but it is out of order. This should not impact the user, but may impact scripts or applications that are parsing the output of the command 	

4.5.6.2.21 show running-config interface

	show running-config interface [lo <loopback_id>] Displays running-config filtered with the specific interfaces.	
Syntax Description	loopback_id	Loopback interface ID. Range: 0-31
Default	N/A	
Configuration Mode	config	
History	3.8.1000	

	3.8.3000	Updates command
Example	<pre>switch (config) # show running-config interface lo 1 interface loopback 1 interface loopback 1 ip address 1.1.10.10/32 primary</pre>	
Related Commands		
Notes		

4.6 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”.

4.6.1 mDNS Commands

4.6.1.1 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
Example	<pre>switch (config) # no ha dns enable</pre>
Related Commands	
Notes	

5 NTP and Clock



Network Time Protocol (NTP) is a networking protocol for clock synchronization between computer systems over variable-latency data networks. NTP is intended to synchronize all participating computers to within a few milliseconds of Coordinated Universal Time (UTC) and is designed to mitigate the effects of variable network latency. NTP can usually maintain time to within tens of milliseconds over the public Internet, and can achieve better than one millisecond accuracy in local area networks under ideal conditions.

5.1 NTP Authenticate

When authentication of incoming NTP packets is enabled, the switch ensures that they come from an authenticated time source before using them for time synchronization on the switch. Authentication keys are created and added to the trusted list.

To add a key to be used for authentication, take the following steps.

1. Create the key.

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

2. Add the key to the trusted list.

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

3. Assign the key to the server/peer.

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

5.2 NTP Authentication Key

An authentication key may be created and used to authenticate incoming NTP packets. For the key to be used, make sure the following is in place.

1. It should be shared with the NTP server/peer sending the NTP packet.
2. It should be added to the trusted list.
3. NTP authenticate should be enabled on the system

5.3 NTP Commands

- [5.1 NTP Authenticate](#)
- [5.2 NTP Authentication Key](#)
- [5.3 NTP Commands](#)
 - [5.3.1 clock set](#)

- [5.3.2 clock timezone](#)
- [5.3.3 ntp](#)
- [5.3.4 ntpdate](#)
- [5.3.5 ntp authenticate](#)
- [5.3.6 ntp authentication-key](#)
- [5.3.7 ntp peer disable](#)
- [5.3.8 ntp peer keyID](#)
- [5.3.9 ntp peer version](#)
- [5.3.10 ntp server disable](#)
- [5.3.11 ntp server keyID](#)
- [5.3.12 ntp server-role disable](#)
- [5.3.13 ntp server trusted-enable](#)
- [5.3.14 ntp server version](#)
- [5.3.15 ntp trusted-key](#)
- [5.3.16 show clock](#)
- [5.3.17 show ntp](#)
- [5.3.18 show ntp configured](#)
- [5.3.19 show ntp keys](#)

5.3.1 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		
Example		
Related Commands	show clock	
Notes	If not specified, the date will be left the same.	

5.3.2 clock timezone

	<p>clock timezone [<zone-word> [<zone-word> [<zone-word>] [<zone-word>]]]</p> <p>no clock timezone</p> <p>Sets the system time zone. The time zone may be specified in one of three ways:</p> <ul style="list-style-type: none"> • 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 <p>The no form of the command resets time zone to its default (GMT).</p>	
Syntax Description	zone-word	Possible forms this could take include: continent, city, continent, country, city, continent, region, country, city, ocean, and/or island.
Default	GMT	
Configuration Mode	config	
History		
Example		
Related Commands	show clock	
Notes		

5.3.3 ntp

	<p>Configures NTP.</p> <p>The no form of the command negates NTP options.</p>	
Syntax Description	disable	Disables NTP.
	enable	Enables NTP.
	peer server	Configures an NTP peer or server node.
	IP address	IPv4 or IPv6 address.
	version <number>	Specifies the NTP version number of this peer. Possible values: 3 or 4
Default	<p>NTP is enabled</p> <p>NTP version number is 4</p>	
Configuration Mode	config	
History		
Example		
Related Commands		
Notes		

5.3.4 ntpdate

	ntpdate <ip-address> Configures the system clock using the specified SNTP server.	
Syntax Description	ip-address	IP address of SNTP server.
Default	N/A	
Configuration Mode	config	
History		
Example		
Related Commands		
Notes	This is a one-time operation and does not cause the clock to be kept in sync on an ongoing basis. It will generate an error if SNTP is enabled since the socket it requires will already be in use.	

5.3.5 ntp authenticate

	ntp authenticate no ntp authenticate Enables NTP authentication. The no form of the command disables NTP authentication.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config	
History		
Example		
Related Commands		
Notes		

5.3.6 ntp authentication-key

	ntp authentication-key <key-id> <encrypt-type> [<password>] no ntp authentication-key <key-id> Enables NTP authentication. The no form of the command disables NTP authentication.	
Syntax Description	key-id	Specifies a key ID, whether existing or a new one to be added. Range: 1-65534
	encrypt-type	Specifies encryption type to use (md5, or sha1)
	password	Password string
Default	Disabled	
Configuration Mode	config	
History		
Example		

Related Commands	
Notes	If a password is not entered, a prompt appears requiring that a password is introduced.

5.3.7 ntp peer disable

	<pre>ntp peer <ip-address> disable no ntp peer <ip-address> disable</pre> <p>Temporarily disables this NTP peer. The no form of the command enables this NTP peer.</p>	
Syntax Description	ip-address	IP address of the peer.
Default	Disabled	
Configuration Mode	config	
History		
Example		
Related Commands		
Notes	<ul style="list-style-type: none"> IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone ID for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7'.) The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen. 	

5.3.8 ntp peer keyID

	<pre>ntp peer <ip-address> keyID <key-id> no ntp peer <ip-address> keyID <key-id></pre> <p>Specifies the KeyID of the NTP peer. The no form of the command removes key ID configuration from the NTP peer.</p>	
Syntax Description	ip-address	IP address of the peer.
	key-id	Range: 1-65534
Default	Disabled	
Configuration Mode	config	
History		
Example		
Related Commands		
Notes	<ul style="list-style-type: none"> IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone ID for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7'.) The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen. 	

5.3.9 ntp peer version

	<pre>ntp peer <ip-address> version <ver-num></pre> <pre>no ntp peer <ip-address> version <ver-num></pre> <p>Specifies the NTP version number of this peer. The no form of the command defaults NTP to version 4.</p>	
Syntax Description	ip-address	IP address of the peer.
	ver-num	NTP version. Possible values: 3 or 4
Default	4	
Configuration Mode	config	
History		
Example		
Related Commands		
Notes	<ul style="list-style-type: none"> IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone ID for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7') The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen. 	

5.3.10 ntp server disable

	<pre>ntp server <ip-address> disable</pre> <pre>no ntp server <ip-address> disable</pre> <p>Temporarily disables this NTP server. The no form of the command enables this NTP server.</p>	
Syntax Description	ip-address	IP address of the peer.
Default	Disabled	
Configuration Mode	config	
History		
Example		
Related Commands		
Notes	<ul style="list-style-type: none"> IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone ID for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7'.) The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen. 	

5.3.11 ntp server keyID

	ntp server <ip-address> keyID <key-id> no ntp server <ip-address> keyID <key-id> Specifies the KeyID of the NTP server. The no form of the command removes key ID configuration from the NTP server.	
Syntax Description	ip-address	IP address of the peer.
	key-id	Range: 1-65534
Default	Disabled	
Configuration Mode	config	
History		
Example		
Related Commands		
Notes	<ul style="list-style-type: none"> IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone ID for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7'.) The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen. 	

5.3.12 ntp server-role disable

	ntp server-role disable no ntp server-role disable	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Configure terminal	
History		
Role	Admin	
Example		
Related Commands	show ntp	
Notes	This command is configurable.	

5.3.13 ntp server trusted-enable

	ntp server <ip-address> trusted-enable no ntp server <ip-address> trusted-enable Trusts this NTP server; if authentication is configured this will additionally force all time updates to only use trusted servers. The no form of the command removes trust from this NTP server.	
Syntax Description	ip-address	IP address of the peer.
Default	N/A	

Configuration Mode	config
History	
Example	
Related Commands	
Notes	<ul style="list-style-type: none"> • IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone ID for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7'.) • The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen. • NTP trusted servers can be used as a mitigation for Sybil attacks which is a vulnerability caused by NTP peers sharing the same NTP key base. This mitigation adds the concept of trusted servers which if enabled in conjunction with NTP authentication ensures that time information will only be obtained from trusted servers.

5.3.14 ntp server version

	<pre>ntp server <ip-address> version <ver-num></pre> <pre>no ntp server <ip-address> version <ver-num></pre> <p>Specifies the NTP version number of this server. The no form of the command defaults NTP to version 4.</p>	
Syntax Description	ip-address	IP address of the peer.
	ver-num	NTP version. Possible values: 3 or 4
Default	4	
Configuration Mode	config	
History		
Example		
Related Commands		
Notes	<ul style="list-style-type: none"> • IP addresses must be in IPv4 format (e.g., '192.168.0.1') or IPv6 format with scope zone ID for IPv6 link-local addresses (e.g., '2001:db8:701f::8f9' or 'fe80::21c:23f:ec1:4fb%7') • The length of a hostname is limited to 255 characters. Each label (node delimited by a dot in the hostname) is limited to 63 characters and may contain letters, numbers and hyphens ('-'), but may not begin with a hyphen 	

5.3.15 ntp trusted-key

	<pre>ntp trusted-key <key(s)></pre> <pre>no ntp trusted-key <key(s)></pre> <p>Adds one or more keys to the trusted key list. The no form of the command removes keys from the trusted key list.</p>	
Syntax Description	key(s)	Range: 1-65534
Default	Disabled	
Configuration Mode	config	

History	
Example	
Related Commands	
Notes	Keys may be separated with commas without any space, or they may be set as a range using a hyphen.

5.3.16 show clock

	<code>show clock</code> Displays the current system time, date and time zone.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	
Example	<pre>Time: 02:48:41 Date: 2018/1/1 Time zone: UTC (Etc/UTC) UTC offset: same as UTC</pre>
Related Commands	
Notes	

5.3.17 show ntp

	<code>show ntp</code> Displays the current NTP settings.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	
Example	
Related Commands	
Notes	

5.3.18 show ntp configured

	<code>show ntp configured</code> Displays NTP configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode

History	
Example	
	<pre> NTP enabled: yes NTP Authentication enabled: no NTP peer 0.us.pool.ntp.org # Hostname peer configuration Resolved as: 45.79.111.114 Enabled: yes NTP version: 4 Key ID: none NTP peer 2.3.1.3 # IP peer configuration Enabled: yes NTP version: 4 Key ID: none NTP server vnc23 # Hostname server configuration Resolved as: 10.7.2.23 Enabled: yes NTP version: 4 Key ID: none Trusted: no NTP server 1.2.3.4 # IP server configuration Enabled: yes NTP version: 4 Key ID: none Trusted: no NTP server idontexist (DNS resolution failed. Reset or reconfigure NTP to try again) Enabled: yes NTP version: 4 Key ID: none Trusted: no </pre>
Related Commands	
Notes	

5.3.19 show ntp keys

	show ntp configured Displays NTP keys.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	
Example	<pre> NTP Key 1 Trusted: yes Encryption Type: MD5 NTP Key 2 Trusted: yes Encryption Type: MD5 NTP Key 3 Trusted: yes Encryption Type: MD5 NTP Key 4 Trusted: yes Encryption Type: md5 </pre>
Related Commands	
Notes	

6 Network Management Interfaces



6.1 SNMP

Simple Network Management Protocol (SNMP), is a network protocol for the management of a network and the monitoring of network devices and their functions. SNMP supports asynchronous event (trap) notifications and queries. MLNX-OS supports:

- SNMP versions v1, v2c and v3
- SNMP trap notifications
- Standard MIBs
- Private MIBs

6.1.1 Standard MIBs

The following table presents the supported textual conventions and conformance MIBs:

MIB	Standard
INET-ADDRESS-MIB	RFC-4001
SNMPV2-CONF	
SNMPV2-TC	RFC 2579
SNMPV2-TM	RFC 3417
SNMP-USM-AES-MIB	RFC 3826
IANA-LANGUAGE-MIB	RFC 2591
IANA-RTPROTO-MIB	RFC 2932
IANAifType-MIB	
IANA-ADDRESS-FAMILY-NUMBERS-MIB	
IGMP-STD-MIB	RFC2933 (See IGMP-STD-MIB Information section)

The following table presents the supported chassis and switch MIBs:

MIB	Standard	Comments
RFC1213-MIB	RFC 1213	
IF-MIB	RFC 2863	ifXTable only supported.
ENTITY-MIB	RFC 4133	
ENTITY-STATE-MIB	RFC 4268	Fan and temperature states
ENTITY-SENSOR-MIB	RFC 3433	<ul style="list-style-type: none">• Port module transmit/receiver power sensors (for 1U systems only)• Fan and temperature sensors

6.1.2 Private MIBs

MIB	Description
MELLANOX-SMI-MIB	Private MIB main structure (no objects)
MELLANOX-PRODUCTS-MIB	List of OID - per managed system (sysObjID)
MELLANOX-IF-VPI-MIB	IfTable extensions
MELLANOX-EFM-MIB	Partially deprecated MIB (based on Mellanox-MIB) Traps definitions and test trap set scalar are supported.
MELLANOX-ENTITY-MIB	Enhances the standard ENTITY-MIB (contains GUID and ASIC revision).
MELLANOX-POWER-CYCLE	Allows rebooting the switch system
MELLANOX-SW-UPDATE-MIB	Allows viewing what SW images are installed, uploading and installing new SW images
MELLANOX-CONFIG-DB	Allows loading, uploading, or deleting configuration files
MELLANOX-ENTITY-STATE-MIB	Extension to support state change traps Note: Currently supported for power supply insertion and extraction only
MELLANOX-XSTP-MIB	Extension to support STP information
MELLANOX-DCB-TRAPS	Extension traps for ETC and PFC
MELLANOX-QOS	Proprietary QoS MIBs
MELLANOX-WJH-MIB	Defines what-just-happened traps

Private MIBs can be downloaded from the [support](#) website.

6.1.3 Proprietary Traps

The following private traps are supported by the MLNX-OS

MELLANOX-EFM-MIB:

Trap	Action Required
asicChipDown	Reboot the system.
asicOverTempReset	Check fans and environmental temperature.
asicOverTemp	Check fans and environmental temperature.
lowPower	Add/connect power supplies.
internalBusError	N/A
procCrash	Generate SysDump and contact support.
cpuUtilHigh	N/A
procUnexpectedExit	Generate SysDump and contact support.
diskSpaceLow	Clean images and sysDump files using the commands “image delete” and “file debug-dump delete”.
systemHealthStatus	Refer to Health Status table.

Trap	Action Required
lowPowerRecover	N/A
insufficientFans	Check Fans and environmental conditions.
insufficientFansRecover	N/A
insufficientPower	Add/connect power supplies, or change power mode using the command “power redundancy mode”.
insufficientPowerRecover	N/A

For additional information refer to MELLANOX-EFM-MIB.

For event-to-MIB mapping, please refer to [“Supported Event Notifications and MIB Mapping”](#).

The only MELLANOX-POWER-CYCLE trap supported is `mellanoxPowerCyclePlannedReload`.

6.1.4 Configuring SNMP

Activate the SNMP server on your switch by running:

```
switch (config) # snmp-server enable
switch (config) # snmp-server enable notify
switch (config) # snmp-server community public ro
switch (config) # snmp-server contact "contact name"
switch (config) # snmp-server host <host IP address> traps version 2c public
switch (config) # snmp-server location "location name"
switch (config) # snmp-server user admin v3 enable
switch (config) # snmp-server user admin v3 prompt auth md5 priv des
```

Community strings are case sensitive.

Modular switches require SNMP timeout configuration on the agent of 60 seconds.

6.1.5 Resetting SNMPv3 Engine ID

Resetting SNMP engine ID is not supported on modular switch systems.

Switch systems shipped with an OS versions older than 3.6.6102 have all had the exact same SNMPv3 engine ID. Going forward, however, all switch systems will ship with a system-specific engine ID.

Upgrading the OS version to 3.6.6102 or higher does not automatically change the current engine ID. That can be done through one of the following methods after performing the software upgrade:

- Changing a switch system’s profile

- Running “reset factory”
- Using the command “snmp-server engineID reset” (for more details, please see the procedure below)

To reset SNMP engine ID using “snmp-server engineID reset”:

Prerequisites:

If any of the following SNMP configurations exist, please delete/disable them and re-enable/reconfigure them only after SNMP engine ID reset is performed:

1. Make sure SNMP is disabled. Run:

```
switch (config) # no snmp-server enable
```

2. Make sure no SNMP trap host is configured. Run:

```
switch (config) # no snmp-server host <ip-address>
```

3. Make sure no SNMP users are configured. Run:

```
switch (config) # no snmp-server user <username> v3
```

Procedure:

1. Check existing engine ID:

```
switch (config) # show snmp engineID
Local SNMP engineID: <current_key>
```

2. Reset existing engine ID:

```
switch (config) # snmp-server engineID reset
```

3. Verify new engine ID:

```
switch (config) # show snmp engineID
Local SNMP engineID: <new_key>
```

6.1.6 Configuring an SNMPv3 User

To configure an SNMPv3 user:

1. Configure the user using the command:

```
switch (config) # snmp-server user [role] v3 prompt auth <hash type> priv <privacy type>
```

Where:

- user role—admin
 - auth type—md5 or sha or sha224 or sha256 or sha384 or sha512
 - priv type—des or aes-128 or 3des or aes-192 or aes-256 or aes-192-cfb or aes-256-cfb
2. Enter authentication password and its confirmation.
 3. Enter privacy password and its confirmation:

```
switch (config) # snmp-server user admin v3 prompt auth md5 priv des
Auth password: *****
Confirm: *****
Privacy password: *****
Confirm: *****
```

To retrieve the system table, run the following SNMP command:

```
snmpwalk -v3 -l authPriv -a MD5 -u admin -A "<Authentication password>" -x DES -X "<privacy password>"
<system ip> SNMPv2-MIB::system
```

6.1.7 Configuring SNMP Notifications (Traps or Informs)

1. Make sure SNMP and SNMP notification are enable. Run:

```
switch (config) # snmp-server enable
switch (config) # snmp-server enable notify
```

2. Configure SNMP host with the desired arguments (IP Address, SNMP version, authentication methods). More than one host can be configured. Each host may have different attributes. Run:

```
switch (config) # snmp-server host 10.134.47.3 traps version 3 user my-username auth sha my-password
```

3. Verify the SNMP host configuration. Run:

```
switch (config) # show snmp host
Notifications enabled:      yes
Default notification community: public
Default notification port:  162

Notification sinks:

 10.134.47.3
   Enabled:                  yes
   Port:                     162 (default)
   Notification type:       SNMP v3 trap
   Username:                 my-username
   Authentication type:     sha
   Privacy type:            aes-128
   Authentication password: (set)
   Privacy password:        (set)
```

4. Configure the desired event to be sent via SNMP. Run:

```
switch (config) # snmp-server notify event interface-up
```

This particular event is used as an example only.

5. Verify the list of traps and informs being sent to out of the system. Run:

```
switch (config) # show snmp events
Events for which traps will be sent:
asic-chip-down: ASIC (Chip) Down
cpu-util-high: CPU utilization has risen too high
disk-space-low: Filesystem free space has fallen too low
health-module-status: Health module Status
insufficient-fans: Insufficient amount of fans in system
insufficient-fans-recover: Insufficient amount of fans in system recovered
insufficient-power: Insufficient power supply
interface-down: An interface's link state has changed to down
interface-up: An interface's link state has changed to up
internal-bus-error: Internal bus (I2C) Error
liveness-failure: A process in the system was detected as hung
low-power: Low power supply
low-power-recover: Low power supply Recover
new_root: local bridge became a root bridge
paging-high: Paging activity has risen too high
```

```
power-redundancy-mismatch: Power redundancy mismatch
process-crash: A process in the system has crashed
process-exit: A process in the system unexpectedly exited
snmp-authtrap: An SNMP v3 request has failed authentication
topology_change: local bridge triggered a topology change
unexpected-shutdown: Unexpected system shutdown
```

To print event notifications to the terminal (SSH or CONSOLE) refer to [“Monitor”](#).

For the SNMPv1 traps or informs, by default, the "agent address" field is set to the IP address of the "mgmt0" interface. In the case that "source interface" is configured to the same VRF which is used for SNMPv1 traps or informs, the IP address of the source interface is used for "agent address" field. In other cases (e.g., if source interface might be configured in some other VRF), "127.0.0.1" is used for the "agent address".

6.1.8 SNMP SET Operations

The OS allows the user to use SET operations via SNMP interface. This is needed to configure a user/ community supporting SET operations.

6.1.8.1 Enabling SNMP SET

To allow SNMP SET operations using SNMPv1/v2:

1. Enable SNMP communities. Run:

```
switch (config) # snmp-server enable communities
```

2. Configure a read-write community. Run:

```
switch (config) # snmp-server community my-community-name rw
```

3. Make sure SNMP communities are enabled (they are enabled by default). Make sure “(DISABLED)” does not appear beside “Read-only communities” / “Read-write communities”.
Run:

```
switch (config) # show snmp
SNMP enabled   : yes
SNMP port      : 161
System contact :
System location:

Read-only communities:
  public

Read-write communities:
  my-community-name

Interface listen enabled: yes

Listen Interfaces:
  Interface: mgmt0

switch (config) # show snmp
No Listen Interfaces.
```

4. Configure this RW community in your MIB browser.

To allow SNMP SET operations using SNMPv3:

1. Create an SNMPv3 user. Run:

```
switch (config) # snmp-server user myuser v3 auth sha <password1> priv aes-128 <password2>
```

It is possible to use other configuration options not specified in the example above. Please refer to the command [“snmp-server user”](#) for more information.

2. Make sure the username is enabled for SET access and has admin capability level. Run:

```
switch (config) # show snmp user
User name: myuser
Enabled overall:      yes
Authentication type:  sha
Privacy type:         aes-128
Authentication password: (set)
Privacy password:    (set)
Require privacy:     yes
SET access:
  Enabled:            yes
  Capability level:   admin
```

The OS supports the OIDs for SET operation listed in the following table which are expanded upon in the following subsections.

	OID Name	OID
MELLANOX-EFM-MIB	sendTestTrapSet	1.3.6.1.4.1.33049.2.1.1.1.6.0
SNMPv2-MIB	sysName	1.3.6.1.2.1.1.5.0
MELLANOX-CONFIG-DB	mellanoxConfigDBCmdExecute	1.3.6.1.4.1.33049.12.1.1.2.3.0
	mellanoxConfigDBCmdFilename	1.3.6.1.4.1.33049.12.1.1.2.2.0
	mellanoxConfigDBCmdStatus	1.3.6.1.4.1.33049.12.1.1.2.4.0
	mellanoxConfigDBCmdStatusString	1.3.6.1.4.1.33049.12.1.1.2.5.0
	mellanoxConfigDBCmdUri	1.3.6.1.4.1.33049.12.1.1.2.1.0
MELLANOX-POWER-CYCLE	mellanoxPowerCycleCmdExecute	1.3.6.1.4.1.33049.10.1.1.2.1.0
	mellanoxPowerCycleCmdStatus	1.3.6.1.4.1.33049.10.1.1.2.2.0
	mellanoxPowerCycleCmdStatusString	1.3.6.1.4.1.33049.10.1.1.2.3.0
MELLANOX-SW-UPDATE	mellanoxSWUpdateCmdSetNext	1.3.6.1.4.1.33049.11.1.1.2.1.0
	mellanoxSWUpdateCmdUri	1.3.6.1.4.1.33049.11.1.1.2.2.0
	mellanoxSWUpdateCmdExecute	1.3.6.1.4.1.33049.11.1.1.2.3.0
	mellanoxSWUpdateCmdStatus	1.3.6.1.4.1.33049.11.1.1.2.4.0
	mellanoxSWUpdateCmdStatusString	1.3.6.1.4.1.33049.11.1.1.2.5.0
	mellanoxSWActivePartition	1.3.6.1.4.1.33049.11.1.1.3.0.0
	mellanoxSWNextBootPartition	1.3.6.1.4.1.33049.11.1.1.4.0.0

6.1.8.2 Sending a Test Trap SET Request

The OS allows the user to use test the notification mechanism via SNMP SET. Sending a SET request with the designated OID triggers a test trap.

Prerequisites:

1. Enable SET operations by following the instructions in [“Enabling SNMP SET”](#).
2. Configure host to which to send SNMP notifications.
3. Set a trap receiver in the MIB browser.

Procedure:

1. Send a SET request to the switch IP with the OID 1.3.6.1.4.1.33049.2.1.1.1.6.0.
2. Make sure the test trap is received by the aforementioned trap receiver (OID: 1.3.6.1.4.1.33049.2.1.2.13).

6.1.8.3 Setting Hostname with SNMP

The OS supports setting system hostname using an SNMP SET request as described in SNMPv2-MIB (sysName, OID: 1.3.6.1.2.1.1.5.0).

The restrictions on setting a hostname via CLI also apply to setting a hostname through SNMP. Refer to the command “hostname” for more information.

6.1.8.4 Power Cycle with SNMP

The OS supports power cycling its systems using an SNMP SET request as described in MELLANOX-POWER-CYCLE MIB.

Power cycle command is issued via the OID `mellanoxPowerCycleCmdExecute`. The following options are available:

- Reload—saves any unsaved configuration and reloads the switch
- Reload discard—reboots the system and discards of any unsaved changes
- Reload force—forces an expedited reload on the system even if it is busy without saving unsaved configuration (equals the CLI command `reload force`)
- Reload slave—reloads the slave management on dual management systems (must be executed from the master management module)

On modular switch systems, it is advised to connect via the BIP to make sure commands are executed from the master management.

6.1.8.5 Changing Configuration with SNMP

The OS supports making configuration changes on its systems using SNMP SET requests. Configuration requests are performed by setting several values (arguments) and then executing a command by setting the value for the relevant operation.

It is possible to set the parameters and execute the commands on the same SNMP request or separate them to several SET operations. Upon executing a command, the values of its arguments remain and can be read using GET commands.

Once a command is executed there may be two types of errors:

- Immediate: This error results in a failure of the SNMP request. This means a critical error in the SNMP request has occurred or that a previous SET request is being executed
- Delayed: The SET request has been accepted by the switch but an error occurred during its execution.

For example, when performing a fetch (download) operation, an immediate error can occur when the given URL is invalid. A delayed error can occur if the download process fails due to network connectivity issues.

The following parameters are arguments are supported:

- Command URI—URI to fetch the configuration file from or upload the file to (for supported URI format please refer to the CLI command “configuration fetch” for more details)
- Config file name—filename to save the configuration file to or to upload to remote location

The following commands are supported:

- BinarySwitchTo—replaces the configuration file with a new binary configuration file. This option fetches the configuration file from the URI provided in the `mellanoxConfigDBCmUri` and switches to that configuration file. This command should be preceded by a `reload` command in order for the new configuration to apply.
- TextApply—fetches a configuration file in human-readable format and applies its configuration upon the current configuration.
- BinaryUpload—uploads a binary format configuration file of the current running configuration or an existing configuration file on the switch to the URI in the `mellanoxConfigDBCmUri` command. The filename parameter indicates what configuration file on the switch to upload.
- TextUpload—uploads a human-readable configuration file of the current running configuration or an existing configuration file on the switch to the URI in the `mellanoxConfigDBCmUri` command. The filename parameter indicates what configuration file on the switch to upload (same as the CLI command `configuration text generate file <filename> upload`).
- ConfigWrite—saves active configuration to a filename on the switch as given in the filename parameter. In case filename is “active”, active configuration is saved to the current saved configuration (same as the CLI command `configuration write`).
- BinaryDelete—deletes a binary based configuration file
- TextDelete—deletes a text based configuration file

6.1.8.6 Upgrading OS Software with SNMP

The OS supports upgrading its software using an SNMP SET request as described in MELLANOX-SW-UPDATE MIB.

The software upgrade command is issued via the OID `mellanoxSWUpdateCmdExecute`. The following options are available:

- Update—fetches the image from a specified URI (equivalent to the command “image fetch” followed by “image install”)
The image to update from is defined by the OID `mellanoxSWUpdateCmdUri`. The restrictions on the URI are identical to what is supported in the CLI command [“image fetch”](#).
- Set-Next—changes the image for the next boot equivalent to the CLI command “image boot”)
The partition from which to boot is defined by the OID `mellanoxSWUpdateCmdSetNext`. The parameters for this OID are as follows:
 - 0—no change
 - 1—partition 1
 - 2—partition 2
 - 3—next partition (default)

Using the OIDs `mellanoxSWUpdateCmdStatus` and `mellanoxSWUpdateCmdStatusString`, you may view the status of the latest operation performed from the aforementioned in either integer values, or human-readable forms, respectively. The integer values presented may be as follows:

- 0—no operation
- 1-100—progress in percentage
- 101—success
- 200—failure

6.1.8.7 IF-MIB and Interface Information

The OS supports displaying information of switch ports, LAG ports, MLAG ports and VLAN interfaces on all systems via SNMP interface. This feature is enabled by default. The interface information is available in the ifTables, ifXTable and mellanoxIfVPITable.

Additionally, traps for interface up/down, and internal link suboptimal speed are enabled. It is possible to enable one or both of these traps.

Interface up/down traps are sent whenever there is a change in the interface's operational state. These traps are suppressed for internal links when the internal link's speed does not match the configured speed of the link (mismatch condition).

6.1.9 Additional Readings and Use Cases

For more information about this feature and its potential applications, please refer to the following community posts:

- [Getting Started with SNMP MIBs](#)
- [HowTo Use SNMP SET](#)

6.2 JSON API

JavaScript Object Notation (JSON) is a machine-to-machine data-interchange format which is supported in the CLI.

The JSON API allows executing CLI commands and receiving outputs in JSON format which can be easily parsed by the calling software.

6.2.1 Authentication

The JSON API protocol runs over HTTP/HTTPS and uses the existing web authentication mechanism.

In order to access the system via HTTP/HTTPS, an HTTP/HTTPS client is needed to send POST requests to the system.

HTTPS access to the web-based management console needs to be enabled using the command “web https enable” to allow POST requests.

The HTTPS client must first be authenticated by sending a POST request to the following URL:

```
https://<ip-address>/admin/launch?script=rh&template=json-request&action=json-login
```


The POST request content should contain the following data (may also be saved as a file) in a JSON format:

```
{
  "username": "<user name>",
  "password": "<user password>"
}
```

After a successful login, a session ID (cookie) is returned to be used for other HTTPS requests in the system.

6.2.1.1 Authentication Example

Before sending JSON HTTPS request, the user must first authenticate.

Create a JSON format file that contains the relevant login credentials. For example, add this content to a file called "post.json":

```
{
  "username": "admin",
  "password": "admin"
}
```

Run the following from your server's shell to create a login session ID in the file: cookiejar.

```
curl -L -X POST -d @post.json -c cookiejar "http://<ip-address>/admin/launch?script=rh&template=json-request&action=json-login"
```

Upon a successful login, you will receive a reply similar to the following:

```
{
  "status": "OK",
  "status_message": "Successfully logged-in"
}
```

The session ID can now be used in all other JSON HTTPS requests to the system.

If authentication fails, the following message is received:

```
{
  "status": "ERROR",
  "status_message": "<Invalid username or password | Please provide username and password>"
}
```

You may also log in and execute commands in the same JSON request. In this case, the JSON file must be in the following format:

```
{
  "username": "<user name>",
  "password": "<user password>",
  "commands | cmd": ["<cli command 1>", "<cli command 2>"] | "<cli command>",
  "execution_type": "sync | async"
}
```

For example:

```
{
  "username": "admin",
  "password": "admin",
  "cmd": "show fan"
}
```

If login is successful, the JSON API response appears. Otherwise, login failure response is presented.

6.2.1.2 Changing Initial Password Through JSON API

This section provides support for changing the default password through JSON API.

Expected Input

- To change the initial password, the payload will be as follows:

```
{
  "username": "admin",
  "password": "admin",
  "initial_admin_password": "admin",
  "initial_monitor_password": "monitor"
}
```

Expected Outputs

- Admin and Monitor passwords cannot be changed because they have already been changed:

```
{
  "status": "ERROR",
  "status_message": " 'admin' password was already set & 'monitor' password was already set"
}
```

- Admin and Monitor passwords were changed successfully:

```
{
  "status_message": " <'admin' password was updated successfully> & <'monitor' password was updated successfully> "
}
```

- Admin and Monitor passwords were not updated:

```
{
  "status": "OK",
  "status_message": "'admin' password was updated successfully & 'monitor' password was updated successfully"
}
```

- One of the passwords of either Admin or Monitor was changed, while the other remained the same:

```
{
  "status": "<ERROR|OK>",
  "status_message": " < Initial password for the 'admin' password was already set | 'admin' password was updated successfully> "
}
```

- When the payload does not have initial passwords, check change-password nodes to see if there is no updated password return in this JSON payload:

```
{
  "status": "ERROR",
  "status_message": "Please set the default password for 'admin' account by using initial password parameters"
}
```

When there is no issue with the login, flow will proceed without needing this step.

6.2.1.3 JSON API Logout

To logout, do the following:

1. Performs a POST operation on URL (the request should contain the session cookie):

```
[switch_ip]/script=rh&template=json-request&action=json-logout
```

2. The switch will remove the session and return the following JSON in the response text (in case of error, content will be relevant to the error):

```
{  
  "status": "OK",  
  "status_message": "Successfully logged-out"  
}
```

3. Make sure there is no cookie. A request with an invalid cookie will respond that the cookie is invalid.

Logout Example

To logout, use the “curl” tool.

```
curl -b cookiejar "http://[switch-ip]/admin/launch?script=rh&template=json-request&action=json-logout"
```

6.2.2 Sending the Request

After successful authentication, the HTTPS client can start sending JSON requests. All requests (POST and GET) should be sent to the following URL:

After the request is handled in the system the HTTPS client receives a JSON response with an indication of the request execution result. If there is data resulting from the request, it is returned as part of the response.

See [“JSON Request Format”](#) for the CLI request format.

See [“JSON Response Format”](#) for the reply format. JSON requests may also be sent using the WebUI. For more information on using the WebUI with JSON, please refer to [“JSON Request Using WebUI”](#).

6.2.3 JSON Request Format

6.2.3.1 JSON Execution Requests

JSON execution requests are HTTPS POST requests that contain CLI commands to be executed in the system.

Execution request can contain a single command or multiple commands to be executed.

Single command execution request format:

```
{  
  "cmd": "<CLI command to execute>"  
}
```

```
}
}
```

Example:

```
{
  "cmd": "show interfaces ethernet 1/1"
}
```

Multiple command execution request format:

```
{
  "commands":["<CLI cmd 1>", "<CLI cmd 2>", ... , <CLI cmd n>"]
}
```

Example:

```
{
  "commands":
  [
    "show interfaces ethernet 1/1",
    "show interfaces ethernet 1/2"
  ]
}
```

In case of a multiple command request, the execution of the commands is done in the order they appear in the execution list. Note that the execution of a multiple command request will be stopped upon first failure. That is, in case the execution of one of the commands fails, none of the remaining commands will be executed.

6.2.3.1.1 Execution Types

Execution requests can be either synchronous (default) or asynchronous.

Synchronous requests will wait for a JSON response from the system. The synchronous request has a defined wait time after which the user will receive a timeout response. The timeout for a synchronous request is configurable by the user and is 30 seconds by default (see the CLI command [“json-gw synchronous-request-timeout”](#)).

Asynchronous requests will return immediately after sending the request with a reply containing a “job_id” key. The user can use the given job ID to later query for request status and execution results. Queries for asynchronous request results are guaranteed to be accessible up to 60 seconds after the request has been completed. After the result has been successfully queried it will be deleted and will no longer be accessible (even if the result is not 60 seconds old).

To specify the execution type, the user needs to add the following key to the JSON execution request:

```
"execution_type": "<async|sync>"
```

Example:

```
{
  "execution_type": "async",
  "cmd": "show interfaces ethernet 1/1"
}
```

6.2.3.2 JSON Query Requests

JSON Query requests are HTTPS GET requests that contain a job ID parameter. Using a query request, the user can get information on the current execution state of an ongoing request or the execution results of a completed request. To send a query request, the user should add the following parameters to the JSON URL:

```
job_id=<job number>
```

Example:

```
https://<switch-ip-address>/admin/launch?script=json&job_id=<job number>
```

See [“JSON Examples”](#) for more examples.

6.2.4 JSON Response Format

Set commands normally do not return any data or output. If a set command does return an output, it will be displayed in the “status_message” field.

6.2.4.1 Single Command Response Format

The HTTPS POST response format structure is a JSON object consisting of 4 name-value pairs as follows:

```
{
  "executed_command": "<CLI command that was executed>",
  "status" = "<OK|ERROR>",
  "status_message" = "<information on the status received>",
  "data" = {the information that was asked for in the request}
}
```

- **executed_command**—the CLI command that was executed in the request
- **status**—the result of the request execution:
 - “OK” if the execution is successful
 - “ERROR” in case of a problem with the execution
- The value type of this key is “string”.
- **data**—a JSON object containing the information requested. Returns an empty string if there is no data.
- **status message**—additional information on the received status. May be empty. The value type of this key is “string”.

Example:

```
{
  "executed_command": "show interfaces ethernet 1/1",
  "status": "OK",
  "status_message": "",
  "data":
  {
    "speed": "40GbE",
    "admin_state": "up"
  }
}
```

See [“JSON Examples”](#) for more examples.

6.2.4.2 Multiple Command Response Format

The HTTPS response format structure is a JSON object consisting of a list of JSON results. Each JSON structure in the list is structured the same as in the single command execution response (see the [previous section](#)).

However, the status field can contain in this case an additional value, “ABORTED”, in case a previous command failed. This status value indicates that the command has not been executed at all in the system.

```
{
  "results": [
    {
      "executed_command": "<...>",
      "status": "<OK|ERROR|ABORTED>",
      "status_message": "<...>",
      "data": {...}
    },
    {
      "executed_command": "<...>",
      "status": "<OK|ERROR|ABORTED>",
      "status_message": "<...>",
      "data": {...}
    },
    ...
    {
      "executed_command": "<...>",
      "status": "<OK|ERROR|ABORTED>",
      "status_message": "<...>",
      "data": {...}
    }
  ]
}
```

Example:

```
{
  "results": [
    {
      "executed_command": "show interfaces ethernet 1/1",
      "status": "OK",
      "status_message": "",
      "data": {"speed": "40GbE", "admin_state": "up"}
    },
    {
      "executed_command": "show interfaces ethernet 1/100",
      "status": "ERROR",
      "status_message": "wrong interfaces name",
      "data": ""
    },
    {
      "executed_command": "show interfaces ethernet 1/2",
      "status": "ABORTED",
      "status_message": "",
      "data": ""
    }
  ]
}
```

See [“JSON Examples”](#) for more examples.

6.2.4.3 Query Response Format

Response to a query request can be of two types. In case the request completes its execution, the response will be similar to the single/multiple command response format, depending on the format of the request, and will display the execution results.

In case the execution is not complete yet, the response format will be similar to the single command response format. However, the status field will contain in this case the value “PENDING” to indicate that the request is still in progress. In addition, the “executed_command” field will contain the current request command being handled by the system.

Example:

```
{
  "executed_command": "show interfaces ethernet 1/1",
  "status": "PENDING",
  "status_message": "",
  "data": ""
}
```

6.2.4.4 Asynchronous Response Format

Response to an asynchronous request is similar to the HTTPS response format of the single command response. However, an additional unique field will be added, “job_id”, containing the job id number for querying the request later. The value of the job_id key is of type string.

Another difference is that the “executed_command” field will be empty.

Example:

```
{
  "executed_command": ""
  "status": "OK"
  "status_message": ""
  "data": ""
  "job_id": "2754930426"
}
```

6.2.5 Supported Commands

- Show commands
- Set commands—all non-interactive CLI set commands are supported

Interactive commands are commands which require user interaction to complete (e.g. type “yes” to confirm). These commands are not supported by the JSON API.

6.2.6 JSON Examples

The following examples use curl (a common tool in Linux systems) to send HTTPS POST requests to the system.

6.2.6.1 Synchronous Execution Request Example

6.2.6.1.1 Single Command

This example sends a request to query the system profile.

Request (save it to a file named req.json):

```
{"cmd": "show system profile"}
```

Send the request:

```
curl -b /tmp/cookie -X POST -d @req.json "https://10.10.10.10/admin/launch?script=json"
```

When the system finishes processing the request, the user will receive a response similar to the following:

```
{
  "status": "OK",
  "executed_command": "show system profile",
  "status_message": "",
  "data": {
    "Profile": "ib",
    "Adaptive Routing": "yes",
    "Number of SWIDs": "1"
  }
}
```

6.2.6.1.2 Multiple Commands

This example sends a request to change an interface description and then queries for its status.

Request (save it to a file named req.json):

```
{"commands": ["interfaces ib 1/1 description test description",
"show interfaces ib 1/1 status"]}
```

Send the request:

```
curl -b /tmp/cookie -X POST -d @req.json "https://10.10.10.10/admin/launch?script=json"
```

When the system finishes processing the request, the user will receive a response similar to the following:

```
{
  "results": [
    {
      "status": "OK",
      "executed_command": "interfaces ib 1/1 description test description",
      "status_message": "",
      "data": ""
    },
    {
      "status": "OK",
      "executed_command": "show interfaces ib 1/1 status",
      "status_message": "",
      "data": {
        "IB1/1": [
          {
            "Description": "test description",
            "Speed": "56.0 Gbps",
            "Logical port state": "Initialize",
            "Physical port state": "LinkUp",
            "Current line rate": "56.0 Gbps",
            "IB Subnet": "infiniband-default"
          }
        ]
      }
    }
  ]
}
```

6.2.6.2 Asynchronous Execution Request Example

This example sends an asynchronous request to change an interface description and then queries for its status.

Request (save it to a file named req.json):


```
{"execution_type": "async",
"commands": ["interfaces ib 1/1 description test description",
"show interfaces ib 1/1 status"]}
```

Send the request:

```
curl -b /tmp/cookie -X POST -d @req.json "https://10.10.10.10/admin/launch?script=json"
```

The system immediately returns a response similar to the following:

```
{
  "executed_command": "",
  "status": "OK",
  "status_message": "",
  "data": "",
  "job_id": "91329386"
}
```

6.2.6.3 Query Request Example

This example sends a request to query for a job ID received from a previous execution request.

The request is an HTTPS GET operation to the JSON URL with the “job_id” parameter.

Send the request:

```
curl -b /tmp/cookie -X GET "https://10.10.10.10/admin/launch?script=json&job_id=91329386"
```

If the system is still processing the request, the user receives a response similar to the following:

```
{
  "executed_command": " interfaces ib 1/1 description test description ",
  "status": "PENDING",
  "status_message": "",
  "data": ""
}
```

If the system is done processing the request, the user receives a response similar to the following:

```
{
  "results": [
    {
      "status": "OK",
      "executed_command": "interfaces ib 1/1 description test description",
      "status_message": "",
      "data": ""
    },
    {
      "status": "OK",
      "executed_command": "show interfaces ib 1/1 status",
      "status_message": "",
      "data": {
        "IB1/1": [
          {
            "Description": "test description",
            "Speed": "fdr",
            "Logical port state": "Initialize",
            "Physical port state": "LinkUp",
            "Current line rate": "56.0 Gbps",
            "IB Subnet": "infiniband-default"
          }
        ]
      }
    }
  ]
}
```

6.2.6.4 Error Response Example

6.2.6.4.1 General Error

This example sends a request with an illegal JSON structure.

Request—without closing bracket “]” (save it to a file named req.json):

```
{"commands": ["interface ib 1/1 description test description",  
"show interfaces ib 1/1 status"]
```

Send the request:

```
curl -b /tmp/cookie -X POST -d @req.json "https://10.10.10.10/admin/launch?script=json"
```

Error response:

```
{  
  "status": "ERROR",  
  "executed_command": "",  
  "status_message": "Handle request failed. Reason:\nIllegal JSON structure found in given JSON data.  
\nExpecting , delimiter: line 1 column 95 (char 94)",  
  "data": ""  
}
```

6.2.6.4.2 Multiple Command Request Failure

This example sends a multiple command request where one command fails.

Request—with a non-existing interface (1/200) (save it to a file named req.json):

```
{  
  "execution_type": "sync",  
  "commands": [  
    "interfaces ib 1/1 speed sdr",  
    "interfaces ib 1/200 speed sdr",  
    "interfaces ib 1/3 speed sdr"]  
}
```

Send the request:

```
curl -b /tmp/cookie -X POST -d @req.json "https://10.10.10.10/admin/launch?script=json"
```

Error response:

```
{  
  "results": [  
    {  
      "status": "OK",  
      "executed_command": "interfaces ib 1/1 speed sdr",  
      "status_message": "",  
      "data": ""  
    },  
    {  
      "status": "ERROR",  
      "executed_command": "interfaces ib 1/200 speed sdr",  
      "status_message": "% 1st interface does not exist",  
      "data": ""  
    }  
  ],  
}
```

```
{
  "status": "ABORTED",
  "executed_command": "interfacesib 1/3 speed sdr",
  "status_message": "",
  "data": ""
}
]
```

6.2.7 JSON Request Using WebUI

The
MLNX-OS

WebUI also allows users to send JSON HTTPS POST and GET requests.

Log into the WebUI, go to the “Setup” tab, and select “JSON API” from the left side menu.

This section is displayed only if JSON API is enabled using the command “json-gw enable”.

6.2.7.1 To Execute a JSON Request

1. Choose “Execute JSON command”.
2. Choose the “execution_type” from the drop down list.
3. In the “commands” field, type the CLI command(s) to execute.
Use the “+” and “-” buttons to add or remove additional commands to the request.
4. Click “Submit”.

The JSON response is then shown in the “JSON Response” box below.

The HTTPS method (HTTPS POST in this instance) and the URL used to send the request will be displayed next to the “HTTPS Method” and “URL” field respectively.

JSON API **Product Documents**

- Interfaces
- HA
- Routing
- Hostname
- DNS
- Login/Logout Messages
- Address Resolution
- IPSec
- Neighbors
- Virtualization
- Virtual Switch Mgmt
- Web
- SNMP
- Email Alerts
- XML gateway
- JSON API
- Logging
- Configurations
- Date and Time
- NTP
- Licensing

JSON Configuration

Enable JSON API

Apply **Cancel**

JSON Commands

Execute JSON command
 Query asynchronous job status

Enter one or more CLI commands to be executed:

```
{
  "execution_type": "sync",
  "commands": [
    "show system profile"
  ]
}
```

Submit **Cancel**

JSON Response

HTTP Method: **POST** URL: **http://[redacted]/admin/launch?script=json**

```
{
  "results": [
    {
      "status": "OK",
      "executed_command": "show system profile",
      "status_message": "",
      "data": {
        "Profile": "ib",
        "Adaptive Routing": "yes",
        "Number of SWIDs": "1"
      }
    }
  ]
}
```

6.2.7.2 To Query an Asynchronous JSON Request

1. Choose “Query asynchronous job status”.
2. Type the job ID in the “Job ID” text box.
3. Press “Query Status”.

The JSON response is then shown in the “JSON Response” box below. The HTTPS method (HTTPS GET in this instance) and the URL used to send the request will be displayed next to the “HTTPS Method” and “URL” field respectively.

JSON API i
Product Documents

- Interfaces
- HA
- Routing
- Hostname
- DNS
- Login/Logout Messages
- Address Resolution
- IPSec
- Neighbors
- Virtualization
- Virtual Switch Mgmt
- Web
- SNMP
- Email Alerts
- XML gateway
- JSON API
- Logging
- Configurations
- Docker
- Date and Time
- NTP
- Licensing

JSON Configuration

Enable JSON API

Apply **Cancel**

JSON Commands

Execute JSON command

Query asynchronous job status

Job ID

Query Status **Cancel**

JSON Response

HTTP Method: URL:

```

{
  "results": [
    {
      "status": "OK",
      "executed_command": "show system profile",
      "status_message": "",
      "data": {
        "Profile": "vpi-single-switch"
      }
    }
  ]
}

```

6.2.8 Additional Reading and Use Cases

For more information about this feature and its potential applications, please refer to the following community post:

- [Getting Started With JSON API](#)

6.3 Network Management Interface Commands



- **6.3.1 SNMP**
 - [6.3.1.1 snmp-server auto-refresh](#)
 - [6.3.1.2 snmp-server cache enable](#)
 - [6.3.1.3 snmp-server community](#)
 - [6.3.1.4 snmp-server contact](#)
 - [6.3.1.5 snmp-server enable](#)
 - [6.3.1.6 snmp-server engineID reset](#)
 - [6.3.1.7 snmp-server enable mult-communities](#)
 - [6.3.1.8 snmp-server enable notify](#)
 - [6.3.1.9 snmp-server enable set-permission](#)
 - [6.3.1.10 snmp-server host disable](#)
 - [6.3.1.11 snmp-server host informs](#)
 - [6.3.1.12 snmp-server host traps](#)
 - [6.3.1.13 snmp-server listen](#)
 - [6.3.1.14 snmp-server notify](#)
 - [6.3.1.15 snmp-server port](#)

- [6.3.1.16 snmp-server user](#)
- [6.3.1.17 show snmp](#)
- [6.3.1.18 show snmp auto-refresh](#)
- [6.3.1.19 show snmp engineID](#)
- [6.3.1.20 show snmp set-permission](#)
- [6.3.1.21 show snmp user](#)
- [6.3.1.22 show interfaces ib internal notification](#)
- [6.3.2 JSON API](#)
 - [6.3.2.1 json-gw enable](#)
 - [6.3.2.2 json-gw synchronous-request-timeout](#)
 - [6.3.2.3 show json-gw](#)

6.3.1 SNMP

6.3.1.1 snmp-server auto-refresh

	snmp-server auto-refresh {enable interval <time>} no snmp-server auto-refresh enable Configures SNMPD refresh settings. The no form of the command disables SNMPD refresh mechanism.	
Syntax Description	enable	Enables SNMPD refresh mechanism.
	interval	Sets SNMPD refresh interval.
	time	Range: 20-500 seconds
Default	Enabled Interval—60 seconds	
Configuration Mode	config	
History	3.2.3000 3.4.1100: Added “time” parameter and updated notes	
Example	switch (config) # snmp-server auto-refresh interval 120	
Related Commands	show snmp	
Notes	<ul style="list-style-type: none"> • When configuring an interval lower than 60 seconds, the following warning message appears asking for confirmation: “Warning: this configuration may increase CPU utilization, Type 'YES' to confirm: YES • When disabling SNMP auto-refresh, information is retrieved no more than once every 60 seconds just like SNMP tables that do not have an auto-refresh mechanism 	

6.3.1.2 snmp-server cache enable

	snmp-server cache enable no snmp-server cache enable Enables SNMP cache if auto-refresh is disabled. The no form of the command disables SNMP cache if auto-refresh is disabled.	
Syntax Description	N/A	
Default	Enabled	

Configuration Mode	config
History	3.7.0000
Example	<code>switch (config) # snmp-server cache enable</code>
Related Commands	<code>show snmp auto-refresh</code> <code>snmp-server auto-refresh enable</code>
Notes	<ul style="list-style-type: none"> • If SNMP auto-refresh is enabled, the value of cache is meaningless • If SNMP cache is disabled, every SNMP request gets updated data

6.3.1.3 snmp-server community

	<code>snmp-server community <community> [ro rw]</code> <code>no snmp-server community <community></code> Sets a community name for either read-only or read-write SNMP requests. The no form of the command sets the community string to default.	
Syntax Description	community	Community name
	ro	Sets the read-only community string
	rw	Sets the read-write community string
Default	Read-only community: "public" Read-write community: ""	
Configuration Mode	config	
History	3.1.0000	
Example	<code>switch (config) # snmp-server community private rw</code>	
Related Commands	<code>show snmp</code>	
Notes	<ul style="list-style-type: none"> • If neither the "ro" or the "rw" parameters are specified, the read-only community is set as the default community • If the read-only community is specified, only queries can be performed • If the read-write community is specified, both queries and sets can be performed 	

6.3.1.4 snmp-server contact

	<code>snmp-server contact <contact-name></code> <code>no snmp-server contact</code> Sets a value for the sysContact variable in MIB-II. The no form of the command resets the parameter to its default value.	
Syntax Description	contact-name	Contact name
Default	""	
Configuration Mode	config	
History	3.1.0000	
Example	<code>switch (config) # snmp-server contact my-name</code>	
Related Commands	<code>show snmp</code>	
Notes		

6.3.1.5 snmp-server enable

	<code>snmp-server enable</code> <code>no snmp-server enable</code> Enables SNMP-related functionality (SNMP engine, and traps). The no form of the command disables the SNMP server.
Syntax Description	N/A
Default	SNMP is enabled by default
Configuration Mode	config
History	3.1.0000
Example	<code>switch (config) # snmp-server enable</code>
Related Commands	<code>show snmp</code>
Notes	

6.3.1.6 snmp-server engineID reset

	<code>snmp-server engineID reset</code> Resets the SNMPv3 engine ID to be node unique.
Syntax Description	N/A
Default	Default engineID is unchanged
Configuration Mode	config
History	3.6.6102
Example	<code>switch (config) # snmp-server engineID reset</code>
Related Commands	<code>show snmp engineID</code>
Notes	Changing system profile or performing “reset factory...” causes the engine ID to change to the new node-unique one.

6.3.1.7 snmp-server enable multi-communities

	<code>snmp-server enable multi-communities</code> <code>no snmp-server enable multi-communities</code> Enables multiple communities to be configured. The no form of the command disables multiple communities to be configured.
Syntax Description	N/A
Default	SNMP server multi-communities are disabled by default
Configuration Mode	config
History	3.1.0000
Example	<code>switch (config) # snmp-server enable multi-communities</code>
Related Commands	<code>show snmp</code>
Notes	

6.3.1.8 snmp-server enable notify

	snmp-server enable notify no snmp-server enable notify Enables sending of SNMP traps and informs from this system. The no form of the command disables sending of SNMP traps and informs from this system.	
Syntax Description	N/A	
Default	SNMP notifies are enabled by default	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # snmp-server enable notify	
Related Commands	show snmp	
Notes	SNMP traps are only sent if there are trap sinks configured with the “snmp-server host...” command, and if these trap sinks are themselves enabled.	

6.3.1.9 snmp-server enable set-permission

	snmp-server enable set-permission <MIB-name> no snmp-server enable set-permission <MIB-name> Allows SNMP SET requests for items in a specified MIB. The no form of the command disallows SNMP SET requests for items in a specified MIB.	
Syntax Description	N/A	
Default	SNMP MIBs are all given permission for SET requests by default	
Configuration Mode	config	
History	3.6.3004	
Example	switch (config) # snmp-server enable set-permission MELLANOX-SW-UPDATE	
Related Commands	show snmp set-permission	
Notes		

6.3.1.10 snmp-server host disable

	snmp-server host <ip-address> disable no snmp-server host <ip-address> [disable] Temporarily disables sending of all notifications to this host. The no form of the commands resumes sending of all notifications to this host.	
Syntax Description	ip-address	IPv4 or IPv6 address
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # snmp-server host 10.10.10.10 disable	

Related Commands	show snmp snmp-server enable
Notes	

6.3.1.11 snmp-server host informs

	<p>snmp-server host <ip-address> informs [<community> port <port> version 2c version 3 {engineID <engineID> user <name> {auth <hash-type> <auth-password> [priv <privacy-type> [<priv-password>]} encrypted auth ... prompt auth ...}]</p> <p>no snmp-server host <ip-address> informs port</p> <p>Send SNMP v2c informs to this host with the default trap community. The no form of the commands removes a host from which SNMP traps should be sent.</p>	
Syntax Description	IP address	IPv4 or IPv6 address.
	community	Specifies trap community string.
	port	Overrides default UDP port for this trap sink.
	version	Specifies the SNMP version of traps to send to this host.
	engineID	Specifies engine ID of this inform sink.
	user	Specifies username for this inform sink.
	auth	Configures SNMPv3 security parameters, specifying passwords in plaintext on the command line (passwords are always stored encrypted).
	hash-type	<ul style="list-style-type: none"> • MD5 • SHA
	auth-password	Plaintext password to use for authentication. If “priv” is not specified the default privacy algorithm is used with the same privacy password as that specified for authentication.
	priv	Specifies SNMPv3 privacy settings for this user.
	privacy-type	<ul style="list-style-type: none"> • aes-128—uses AES-128 encryption for privacy • des—uses DES encryption for privacy
	priv-password	Plaintext password to use for privacy. If not specified, then auth-password is used.
	encrypted	Configure SNMPv3 security parameters specifying passwords in encrypted form.
	prompt	Configure SNMPv3 security parameters specifying passwords securely in follow-up prompts rather than on the command line.
Default	community—public UDP port—162 version—3	
Configuration Mode	config	
History	3.2.1050	
Example	<pre>switch (config) # snmp-server host 1.1.1.1 informs version 3 engineID 0x800041da04643265363932653432303135 user test auth md5 password priv aes-128 password</pre>	
Related Commands	show snmp snmp-server enable snmp-server host informs version 3	

Notes	
-------	--

6.3.1.12 snmp-server host traps

	<p>snmp-server host <ip-address> traps [<community> port <port> version {1 2c} version 3 {user <name> {auth <hash-type> <auth-password> [priv <privacy-type> [<priv-password>]}] encrypted auth ... prompt auth ...}}]</p> <p>no snmp-server host <ip-address> traps port</p> <p>Send SNMP v2c traps to this host with the default trap community. The no form of the commands removes a host from which SNMP traps should be sent.</p>	
Syntax Description	ip-address	IPv4 or IPv6 address.
	community	Specifies trap community string.
	port	Overrides default UDP port for this trap sink.
	version	Specifies the SNMP version of traps to send to this host.
	user	Specifies username for this inform sink.
	auth	Configures SNMPv3 security parameters, specifying passwords in plaintext on the command line (passwords are always stored encrypted).
	hash-type	<ul style="list-style-type: none"> • MD5 • SHA
	auth-password	Plaintext password to use for authentication. If “priv” is not specified the default privacy algorithm is used with the same privacy password as that specified for authentication.
	priv	Specifies SNMPv3 privacy settings for this user.
	privacy-type	<ul style="list-style-type: none"> • aes-128—uses AES-128 encryption for privacy • des—uses DES encryption for privacy
	priv-password	Plaintext password to use for privacy. If not specified, then auth-password is used.
	encrypted	Configure SNMPv3 security parameters, specifying passwords in encrypted form.
	prompt	Configure SNMPv3 security parameters, specifying passwords securely in follow-up prompts, rather than on the command line.
Default	community—public UDP port—162 version—3	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # snmp-server host 1.1.1.1 informs version 3 user test auth md5 password priv aes-128 password</pre>	
Related Commands	show snmp snmp-server enable snmp-server host informs version 3	
Notes		

6.3.1.13 snmp-server listen

	<code>snmp-server listen {enable interface <ifName>}</code> <code>no snmp-server listen {enable interface <ifName>}</code> Configures SNMP server interface access restrictions. The no form of the command disables the listen interface restricted list for SNMP server.	
Syntax Description	enable	Enables SNMP interface restrictions on access to this system
	ifName	Adds an interface to the “listen” list for SNMP server. For example: “mgmt0”, “mgmt1”
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # snmp listen enable</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.	

6.3.1.14 snmp-server notify

	<code>snmp-server notify {community <community> event <event name> port <port> send-test}</code> <code>no snmp-server notify {community event <event name> port}</code> Configures SNMP notifications (traps and informs). The no form of the commands negate the SNMP notifications.	
Syntax Description	community	Sets the default community for traps sent to hosts which do not have a custom community string set
	event	Specifies which events will be sent as traps
	port	Sets the default port to which traps are sent
	send-test	Sends a test trap
Default	All informs and traps are enabled <code>community—public</code> <code>UDP port—162</code>	
Configuration Mode	config	
History	3.1.0000 3.2.1050: Changed traps to notify	
Example	<pre>switch (config) # snmp-server community public</pre>	
Related Commands	<code>show snmp</code> <code>show snmp events</code>	
Notes	<ul style="list-style-type: none"> This setting is only meaningful if traps are enabled, though the list of hosts may still be edited if traps are disabled Refer to Mellanox MIB file for the list of supported traps 	

6.3.1.15 snmp-server port

	snmp-server port <port> no snmp-server port Sets the UDP listening port for the SNMP agent. The no form of the command resets the parameter to its default value.	
Syntax Description	port	UDP port
Default	161	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # snmp-server port 1000	
Related Commands	show snmp	
Notes		

6.3.1.16 snmp-server user

	snmp-server user {admin <username>} v3 {[encrypted] auth <hash-type> <password> [priv <privacy-type> [<password>]] capability <cap> enable <sets> prompt auth <hash-type> [priv <privacy-type>] require-privacy} no snmp-server user {admin <username>} v3 {[encrypted] auth <hash-type> <password> [priv <privacy-type> [<password>]] capability <cap> enable <sets> prompt auth <hash-type> [priv <privacy-type>]} Specifies an existing username, or a new one to be added. The no form of the command disables access via SNMP v3 for the specified user.	
Syntax Description	v3	Configures SNMPv3 users.
	auth	Configures SNMPv3 security parameters, specifying passwords in plaintext on the command line (note: passwords are always stored encrypted). Available hash-type options are: <md5 sha sha224 sha256 sha384 sha512>.
	capability	Sets capability level for SET requests.
	enable	Enables SNMPv3 access for this user.
	encrypted	Configures SNMPv3 security parameters, specifying passwords in encrypted form.
	prompt	Configures SNMPv3 security parameters, specifying passwords securely in follow-up prompts, rather than on the command line.
	require-privacy	Requires privacy (encryption) for requests from this user.
	priv	Configures SNMPv3 security parameters, specifying which protocol to use for traffic encryption. Available priv-type options: <des 3des aes-128 aes-192 aes-256>.
Default	No SNMP v3 users defined	
Configuration Mode	config	
History	3.1.0000 3.7.0000 3.8.1000: Syntax updated	

Example	switch (config) # snmp-server user admin v3 enable
Related Commands	show snmp user
Notes	<ul style="list-style-type: none"> • The username chosen here may be anything that is valid as a local UNIX username (alphanumeric, plus '-', '_', and '.'), but these usernames are unrelated to, and independent of, local user accounts. That is, they need not have the same capability level as a local user account of the same name. Note that these usernames should not be longer than 31 characters, or they will not work. • The hash algorithm specified is used both to create digests of the authentication and privacy passwords for storage in configuration, and also in HMAC form for the authentication protocol itself • There are three variants of the command, which branch out after the “v3” keyword. If “auth” is used next, the passwords are specified in plaintext on the command line. If “encrypted” is used next, the passwords are specified encrypted (hashed) on the command line. If “prompt-pass” is used, the passwords are not specified on the command line the user is prompted for them when the command is executing. If “priv” is not specified, only the auth password is prompted for. If “priv” is specified, the privacy password is prompted for; entering an empty string for this prompt will result in using the same password specified for authentication. • AES privacy type encryption using the newest algorithm, which means we use aes-blumenthal. For more information see http://www.snmp.com/eso/esoConsortiumMIB.txt. • No more than 30 SNMPv3 users are allowed in the database

6.3.1.17 show snmp

	show snmp [events host] Displays SNMP-server configuration and status.	
Syntax Description	events	SNMP events
	host	List of notification sinks
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000 3.6.8008—Updated example	
Example	<pre>switch (config) # show snmp SNMP enabled : no SNMP port : 161 System contact : Test System location : Boston Read-only communities: public Read-write communities: good Interface listen enabled: yes Listen Interfaces: Interface: mgmt0</pre>	
Related Commands	show snmp	
Notes		

6.3.1.18 show snmp auto-refresh

	<code>show snmp auto-refresh</code> Displays SNMPD refresh mechanism status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.6.6000: Updated example 3.7.0000: Updated example
Example	<pre>switch (config) # show snmp auto-refresh SNMP auto refresh: Auto-refresh enabled: yes Refresh interval (sec): 60 Cache enabled: yes Auto-Refreshed tables: ifTable ifXTable mellanoxIfVPItable</pre>
Related Commands	<code>snmp-server auto-refresh</code>
Notes	

6.3.1.19 show snmp engineID

	<code>show snmp engineID</code> Displays SNMPv3 engine ID key.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.6102
Example	<pre>switch (config) # show snmp engineID Local SNMP engineID: 0x80004f4db1dd435e80accf4a4d4d3031</pre>
Related Commands	<code>snmp-server engineID</code>
Notes	

6.3.1.20 show snmp set-permission

	<code>show snmp set-permission</code> Displays SNMP SET permission settings.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.3004

Example	<pre>switch (config) # show snmp set-permission ----- MIB Name Set Enable ----- MELLANOX-CONFIG-DB-MIB yes MELLANOX-EFM-MIB yes MELLANOX-POWER-CYCLE yes MELLANOX-SW-UPDATE no RFC1213-MIB no</pre>
Related Commands	snmp-server enable set-permission
Notes	

6.3.1.21 show snmp user

	<pre>show snmp user Displays SNMP user information.</pre>
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	<p>3.1.0000</p> <p>3.6.8008: Updated example</p>
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</pre>
Related Commands	show snmp
Notes	

6.3.1.22 show interfaces ib internal notification

	<pre>show interfaces ib internal notification Displays information about internal links notification.</pre>
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	<p>3.3.4318</p> <p>3.4.3000</p> <p>Updated example</p>
Example	<pre>switch (config) # show interfaces ib internal notification ===== Internal links information ===== State change enabled : yes Speed mismatch enabled : yes Periodic notifications : 6 (hours)</pre>

Related Commands	interfaces ib internal notification
Notes	

6.3.2 JSON API

6.3.2.1 json-gw enable

	json-gw enable no json-gw enable Enables the JSON API. The no form of the command disables the JSON API.	
Syntax Description	N/A	
Default	JSON API is enabled	
Configuration Mode	config	
History	3.6.3004	
Example	switch (config) # json-gw enable	
Related Commands	show json-gw	
Notes		

6.3.2.2 json-gw synchronous-request-timeout

	json-gw synchronous-request-timeout <timeout-value> no json-gw synchronous-request-timeout Defines a timeout value for synchronous JSON requests (in seconds). The no form of the command returns the timeout value to its default.	
Syntax Description	timeout-value	Define a timeout value for synchronous JSON requests Range: 0-4294967295
Default	JSON API is enabled	
Configuration Mode	config	
History	3.6.3004	
Example	switch (config) # json-gw synchronous-request-timeout 100	
Related Commands	show json-gw	
Notes		

6.3.2.3 show json-gw

	show json-gw Displays the JSON API setting.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	

History	3.6.3004 3.6.4000: Updated example
Example	switch (config) # show json-gw JSON Gateway enabled: yes Synchronous request timeout: 30 JSON API version: 1.0
Related Commands	json-gw enable json-gw synchronous-request-timeout <time out value>
Notes	

7 Virtualization



The Docker feature is not supported in director systems.

MLNX-OS

allows the user to run their own applications on a Linux docker image embedded in the switch software. The container is a pure application sandbox with resource isolation of both memory and compute from the system code/NOS.

Docker container implementation in the OS enhances its VM support to provide a new set of capabilities:

- Network traffic access

Docker containers are implemented in the OS in the same name-space as the network devices allowing the software to send and receive packets from the switch ports by opening a standard Linux socket over the network devices and using an IP address assigned to the device via the legacy management interface (e.g., JSON over HTTP).

It is recommended to assign a unique port number to the Linux socket to prevent ambiguity of applications between the container and the OS.

- Calling the SDK interfaces

Applications running in the docker container are able to implement a set of tools pertaining only to the container such as telemetry features within the network devices. By calling the switch SDK APIs, it can also read data that is not exposed in the OS user interface, or register to receive events that occur in the system (e.g., port up/down).

The container implementation does not limit the container developer from calling the SDK to set parameters. However this is strongly discouraged as it may cause unexpected system behavior where the OS and the container application manage the same resources.

- Query the Linux tables provisioned by OS such as neighbor cache, routing tables, L3 interfaces attributes etc.

7.1 Limiting the Container's Resources

It is possible to configure multiple containers in dockers, however, they would compete for the same memory and compute resources allocated by the switch software (varies for different systems). To ensure system stability and that no random process is killed to free up memory, it is strongly recommended that all resource configurations done in the container utilize OS user interfaces such as JSON/SNMP and take advantage of the internal loopback interface.

7.1.1 Memory Resources Allocation Protocol

The Linux docker supports a hard limit to control memory resource allocation which limits the container to a given amount of user/system memory.

To set the amount of memory allocated to the container, run the following command:

```
switch (config) # docker start imagename latestver containername init memory 25 label newlabel privileged sdk network docker usb-mount
```

7.1.2 CPU Resource Allocation Protocol

Containers have unrestricted access to the host machine's CPU cycles but it is possible to set a number of constraints to limit the containers' access.

To set up limitations or regulate the containers access to CPU resources, run the following command:

```
docker start imagename latestver containername init cpus 0.2 label new_label privileged sdk network
```

7.2 Upgrade Ramifications

7.2.1 Changing Docker Storage Driver

As a result of the upgrade, the docker's storage driver changes, which may cause a few additional changes:

- The containers and docker images become inaccessible to the user (the docker process will not run)
- The user can reach their old containers after a rollback procedure
- The "no docker" command erases all containers and images, including those that were reachable after rollback. Rolling back after running the "no docker" command may result in failure to create configured containers from unknown images.
- The user is advised to execute the "no docker" command at some point in order to clear unused disk space
- It is possible to reload the Docker images after upgrade with the command: `docker load <image_name>_<image_version>.img.gz`
- The images are presented with tab-tab after "docker load " (in cli)
- It is also possible to load the images after rollback after "no docker" was execute. That means that containers can be restarted after upgrade/rollback if their images are loaded (with "docker load").

It is possible to move containers from the current version to the updated one by executing the following steps:

Before upgrade:

1. Save the container as an image—run the command: “docker commit <container_name> <new_image_name> <new_image_version>”. For example: docker commit my_name my_image my_version. You can see the new image by running: “show docker images”.
2. Save the image—run the command: “docker save <image_name> <image_version> <file_name-optional>”. For example: docker save my_image my_version.
3. Upload the image—save the image to a local repository by running: “image upload <image_file_name> <destination_path>”. For example: image upload my_image_my_version.img.gz scp://username:password@fit150/auto/my_dir. The <image_file_name> is presented after clicking tab-tab.

After upgrade:

1. Start docker—run the “no docker shutdown” command.
2. Fetch the restored image—run the “image fetch <file_name>” command. For example: image fetch scp://username:password@fit150/auto/my_dir/my_image_my_version.img.gz
3. Load the image—run the “docker load <image_file_name>” command. For example: docker load my_image_my_version.img.gz
4. Start a container with the defined image—now that the image with all the content from the container is available in the new environment, start a container with this image. Run the command: “docker start <image_name> <image version> <docker_name> <starting_point> | privileged | label | memory | cpus | usb-mount”. For example: docker start my_image my_version new_container now

After an upgrade operation there is a need to rerun copy-sdk command (in case in use).

7.3 Docker Containers Commands

7.3.1 docker

	docker [logging-level <log-level>] no docker Enables dockers then enters docker configuration context. The no form of the command disables dockers, removes configuration, and deletes all containers and docker images.
Syntax Description	N/A • log-level—logging-level for docker. Possible levels: debug error, fatal info, warn
Default	N/A
Configuration Mode	config
History	3.6.2940 3.9.2300—Added log-level option
Example	switch (config) # docker
Related Commands	
Notes	

7.3.2 docker login

	docker login <username> <cleartext password> [server <server address>] Logs in to remote docker repositories.	
Syntax Description	username	Username
	cleartext password	There are 2 options to enter password using the above command: 1. In command—cleartext 2. Using interactive shell—entering all needed input except the password will prompt the user to provide a password which will not be visible while typing. (masked by *)
	server	The "server" field is not mandatory. In case it is not present, the docker will try to login into docker hub repository.
Default	N/A	
Configuration Mode	config	
History	3.9.1600	
Example	switch (config) # docker login abcd 1234	
Related Commands	show docker login	
Notes		

7.3.3 docker logout

	docker logout [server <server address>] Logs out from remote server.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config	
History	3.9.1600	
Example	switch (config) # docker logout	
Related Commands		
Notes	<ul style="list-style-type: none"> There is no need to provide username as only a single user can be connected to a specific server in any given time 	

7.3.4 commit

	commit <container-name> <image-name> <image-version> Creates a new image from a running container.	
Syntax Description	container-name	Name of the running container to commit (limited to 180 characters)
	image-name	Name of the new image to be created

	image-version	Version of the new image to be created
Default	N/A	
Configuration Mode	config docker	
History	3.6.2940 3.6.8008: Added new character limitation for container-name	
Example	<code>switch (config docker) # commit mycontainer test latest</code>	
Related Commands		
Notes		

7.3.5 copy-sdk

	copy-sdk The command provides access to the switch SDK APIs giving applications running on docker access to the switch hardware.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config docker	
History	3.6.4110 3.8.1000: Updated notes 3.8.2100: Updated notes	
Example	<code>switch (config docker) # copy-sdk</code>	
Related Commands		
Notes	<ul style="list-style-type: none"> • Copying SDK files to a USB mounted folder is not allowed • After an upgrade operation there is a need to rerun copy-sdk command (in case in use). 	

7.3.6 remove image

	remove image <image-name> <image-version> Removes an image from the Linux docker service.	
Syntax Description	image-name	Name of the new image to be deleted
	image-version	Version of the new image to be deleted
Default	N/A	
Configuration Mode	config docker	
History	3.6.3520 3.6.2940	
Example	<code>switch (config docker) # remove image test latest</code>	
Related Commands	docker	
Notes		

7.3.7 exec

	<code>exec <container-name> <program-executable></code> Executes a program within a running container.	
Syntax Description	<code>container-name</code> Name of the running container to commit (limited to 180 characters)	
	<code>program-executable</code> Linux command	
Default	N/A	
Configuration Mode	config docker	
History	3.6.3520 3.6.2940	
Example	<pre>switch (config docker) # exec mycontainer "ls -la"</pre>	
Related Commands	docker	
Notes		

7.3.8 label

	<code>label <label name></code> <code>no label <label name></code> Creates a label which can be used as a shared storage between containers. The no form of the command removes the label.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config docker	
History	3.6.4110	
Example	<pre>switch (config docker) # label new_label</pre>	
Related Commands		
Notes		

7.3.9 load

	<code>load <image-name></code> Loads an image from a TAR archive.	
Syntax Description	<code>image-name</code> Name of the TAR image to be loaded	
Default	N/A	
Configuration Mode	config docker	
History	3.6.2940	
Example	<pre>switch (config docker) # load test</pre>	
Related Commands	docker	
Notes		

7.3.10 pull

	pull <image-name>[:<version>] Pulls a docker image from a docker repository.	
Syntax Description	image-name	Image name Format: Name:Version If only "Name" is provided, "version" defaults to latest
Default	N/A	
Configuration Mode	config docker	
History	3.6.2940	
Example	<pre>switch (config docker) # pull test Using default tag: latest latest: Pulling from library/test 45a2e645736c: Pull complete Digest: sha256:c577af3197aacddf79c5a204cd7f493c8e07ffbce7f88f7600bf19c688c38799 Status: Downloaded newer image for test:latest</pre>	
Related Commands	docker	
Notes		

7.3.11 save

	save <image-name> <image-version> <filename> Saves an image to a TAR archive.	
Syntax Description	image-name	Image name
	image-version	Image version
	filename	Name of the file in which to save the image
Default	N/A	
Configuration Mode	config docker	
History	3.6.2940 3.6.8008: Updated command syntax	
Example	<pre>switch (config docker) # save busybox latest my_image Saving and compressing image: busybox version: latest this could take a while... switch (config docker) #</pre>	
Related Commands	docker docker load	
Notes	After the file is created, the filename gets appended a *.gz suffix.	

7.3.12 shutdown

	<p>shutdown no shutdown</p> <p>Stops all docker containers, and deletes all non-auto containers. The no form of the command enables the docker Linux service and runs all configured auto-start containers</p>
Syntax Description	N/A
Default	N/A
Configuration Mode	config docker
History	3.6.2940
Example	<code>switch (config docker) # no shutdown</code>
Related Commands	docker
Notes	

7.3.13 start

	<p>start <image-name> <image-version> <container-name> <starting-point> [privileged {network sdk}] [cpus <max-cpu-resources>] [memory <max-memory>] [usb-mount] [host-trust [user <username>]] [logging-facility <logging-facility-level>] [user-env <env-string>]</p> <p>no start <container-name></p> <p>Starts a new container from an image. The no form of the command stops a running docker container.</p>	
Syntax Description	image-name	Name of the new image to start.
	image-version	Version of the image to start.
	container-name	Name of the running container to commit (limited to 180 characters).
	privileged	<ul style="list-style-type: none"> network—adds network privileges to the container (--privilege flag) sdk—adds required mounts to use the switch SDK from the container
	starting-point	<ul style="list-style-type: none"> init—persistent, start the container after boot, when system initialization is done data-path-ready—persistent, start the container after boot, when data-path is ready to be configured now—start the container now, this is not persistent now-and-data-path-ready—starts the container now and after boot, when data-path is ready to be configured now-and-init—starts the container now and after boot, when system configuration is done
	cpus	Sets how much of the available CPU resources a container can use (e.g., “cpus 1.5” guarantees at most one and a half of the available CPUs for the container).
	memory	Sets the maximum amount of memory the container can use in MB. The minimum amount of memory to configure is 4MB.
	usb-mount	Enables USB mount to the docker container.

	host-trust	Allows SSH operation from within the container to localhost without the need to supply password.
	logging-facility-level	Available Parameters: auth, authpriv daemon, ftp, kern, local0, local1, local2, local3, local4, local5, local6, local7, lpr, mail, news, syslog, user, uucp
	env-string	Up to 16 user-defined environment variables. User-defined environment variable are separated by a comma (e.g., key1=value1,key2=value2)
Default	N/A	
Configuration Mode	config docker	
History	<p>3.6.2940 3.6.3520: Added “privileged” parameter 3.6.8008: Added the options “now-and-data-path-ready” and “now-and-init”, new character limitation for container-name, and updated the description of the parameter “memory”</p> <p>3.8.1000; Updated syntax description 3.9.2000: Added host-trust option which adds support for SSH operation from within the container to localhost without the need to supply password (when activating host-trust without supplying user, user admin will be used). 3.9.2300: Added logging-facility and user-env options</p>	
Example	<pre>switch (config docker) # start centos latest test now Starting docker container. Please wait (this can take a minute)... switch (config) # docker start imagename latestver containername init cpus 0.2 memory 25</pre>	
Related Commands	docker	
Notes	<ul style="list-style-type: none"> • The no form of the command removes the container if it is not persistent. • If trust is set and username is not given, user admin will be chosen by default. 	

7.3.14 image upload

	image upload <filename> <upload_url> Uploads an image file to a remote host.	
Syntax Description	filename	Name of file
	upload_url	FTP, TFTP, SCP and SFTP are supported (e.g., scp://username[:password]@hostname-or-ip/path/filename)
Default	N/A	
Configuration Mode	config	
History	3.6.2940	
Example	<pre>switch (config) # image upload centos.img.gz scp:// username:password@192.168.10.125/var/www/html/<image_name></pre>	
Related Commands		
Notes		

7.3.15 file image upload

	file image upload <filename> <upload_url> Uploads a file to a remote host.	
Syntax Description	filename	Name of file
	upload_url	FTP, TFTP, SCP and SFTP are supported (e.g., scp://username[:password]@hostname/path/filename)
Default	N/A	
Configuration Mode	config	
History	3.6.2940	
Example	switch (config) # file image upload centos.img.gz scp:// username:password@192.168.10.125/var/www/html/<image_name>	
Related Commands		
Notes		

7.3.16 show docker

	show docker Displays docker running state.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.9.2000	
Example	switch (config) # show docker Docker's state: started Docker log-level: warn	
Related Commands		
Notes		

7.3.17 show docker containers

	show docker containers <container_name> Displays set parameters on containers already running, and containers planned to run in the future.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.8008 3.8.1000: Updated example 3.9.2000: Updated example, adding host-trust option 3.9.2300: Updated example, adding "user-defined variables" and "log-facility" fields	

<p>Example</p>	<pre> switch (config) # show docker containers cont_example: image : busybox version : latest status : running start point : data-path-ready cpu limit : 0.2 memory limit: 10m labels : - privileges : network, sdk usb mount : enabled host trust : admin log-facility: kern user-defined variables: name1: value1 name2: value2 another_container: image : busybox version : latest status : - start point : init cpu limit : 0.2 memory limit: 10m labels : my_label privileges : network, sdk usb mount : disabled host trust : admin log-facility: kern user-defined variables: name1: value1 name2: value2 OS_SYSTEM_TYPE : MSB7700 OS_VERSION : 3.9.2300 OS_DOCKERD_VRF_CONTEXT : vrf-default OS_DOCKERD_LINUX_VRF_CONTEXT: vrf_vrf-default switch (config) # show docker containers cont_example cont_example: image : busybox version : latest status : running start point : data-path-ready cpu limit : 0.2 memory limit: 10m labels : - privileges : network, sdk usb mount : enabled host trust : admin log-facility: kern user-defined variables: name1: value1 name2: value2 OS_SYSTEM_TYPE : MSB7700 OS_VERSION : 3.9.2300 OS_DOCKERD_VRF_CONTEXT : vrf-default OS_DOCKERD_LINUX_VRF_CONTEXT: vrf_vrf-default </pre>
<p>Related Commands</p>	

Notes	<ul style="list-style-type: none"> • If a container is already started, the status field displays its current status • If a container is configured to run on the next boot, the start point field displays when it will start • If there is a mismatch between the configuration of a running container and its next-boot configuration, two entries for the container are shown with both of the configurations • For running containers, environment variables that are automatically passed to docker container are revealed (i.e., OS_SYSTEM_TYPE, OS_VERSION, OS_DOCKERD_VRF_CONTEXT, OS_DOCKERD_LINUX_VRF_CONTEXT) • If no user-defined variables were configured, "user-defined variables" field is hidden
-------	---

7.3.18 show docker images

	show docker images Display docker images.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.3520 3.6.2940: Updated example
Example	
<pre>switch (config) # show docker images ----- Image Version Created Size ----- ubuntu latest Less than a secon 117MB d ago ubuntu-sdk v1 41 seconds ago 215MB</pre>	
Related Commands	
Notes	

7.3.19 show docker ps

	show docker ps Display docker containers.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.3520 3.6.2940: Updated example
Example	
<pre>switch (config) # show docker ps ----- Container Image:Version Created Status ----- my_ubuntu_app ubuntu:latest 56 seconds ago Up 50 seconds</pre>	
Related Commands	

Notes	This command is available only after Linux dockers are enabled (“no dockers shutdown”)
-------	--

7.3.20 show docker labels

	show docker labels Displays docker labels.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.4110
Example	<pre>switch (config) # show docker labels Storage label : label_name1 configured containers list : cont_name2 active containers list : cont_name1 Storage label : label_name2</pre>
Related Commands	
Notes	

7.3.21 show docker login

	show docker login Displays docker login.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.9.1600
Example	<pre>switch (config) # show docker login Servers: https://index.docker.io/v1/ nvcr.io</pre>
Related Commands	docker login
Notes	

7.3.22 show docker stats

	show docker stats [<name>] Displays Linux docker statistics.
Syntax Description	name Docker whose stats to display
Default	N/A
Configuration Mode	Any command mode

History	3.6.8008 2.9.2300: Added example
Example	<pre>switch (config) # show docker stats ----- Container CPU % Memory Memory Memory % Block Block Pids Usage Limit OUT ----- container1 0.00% 952K 1000M 0.09% 0B 0B 1</pre>
Related Commands	
Notes	This command is available only after Linux dockers are enabled (“no dockers shutdown”)

8 Telemetry, Monitoring, and Debuggability

- [Logging](#)
- [Link Diagnostic Per Port](#)
- [Event Notifications](#)
- [Buffer Histograms Monitoring](#)
- [Statistics and Alarms](#)
- [Management Information Bases \(MIBs\)](#)

8.1 Logging



8.1.1 Monitor

To print logging events to the terminal, set the modules or events you wish to print to the terminal. For example, run: o-

```
switch (config) # logging monitor events notice
switch (config) # logging monitor sx-sdk warning
```

These commands print system events in severity “notice”, and “sx-sdk” module notifications in severity “warning” to the screen. For example, in case of interface-down event, the following gets printed to the screen:

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

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

8.1.2 Remote Logging

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

1. Set remote syslog server.

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

2. (Optional) Set the destination port of the remote host.

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

3. (Optional) Filter log messages according to an input regex.

```
switch (config) # logging <IP address/hostname> filter <"include"/"exclude"> <regex>
```

4. Set the minimum severity of the log level to info.

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

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

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

8.1.3 Logging Protocol

A feature that provides the ability to choose the protocol to use for sending syslog messages to a remote host: UDP (default) or TCP. See "[logging protocol](#)" command.

8.1.4 Logging Commands

8.1.4.1 logging

	<pre>logging <IPv4 address/IPv6 address/hostname></pre> <pre>no logging <IPv4 address/IPv6 address/hostname></pre> <p>Sends log messages to the remote host specified by its IP or hostname. The no form of the command stops sending log messages to the remote host specified by its IP or hostname.</p>	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config	
History	3.1.1000	
Role	admin	
Example	<pre>switch (config) # logging 1.1.1.1</pre> <pre>switch (config) # no logging 1.1.1.1</pre>	
Related Commands		
Notes	This command is configurable. If "configuration write" is executed, the remote host will still receive messages after reload.	

8.1.4.2 logging port

	<pre>logging <syslog IPv4 address/IPv6 address/hostname> port <destination-port></pre> <pre>no logging <syslog IPv4 address/IPv6 address/hostname> port</pre> <p>Configures remote server destination port for log messages. The no form of the command resets the remote log port to its default value.</p>	
Syntax Description	destination-port	Range: 1-65535
	Hostname	Max 64 characters
Default	514 (UDP)	
Configuration Mode	config	
History	3.6.2002 3.8.1000—Updated command syntax	

Example	switch (config) # logging 10.0.0.1 port 105
Related Commands	logging <syslog IPv4 address/IPv6 address/hostname> trap
Notes	

8.1.4.3 logging trap

	<p>logging <syslog IPv4 address/IPv6 address/hostname> [trap {<log-level> override class <class> priority <log-level>}]</p> <p>no logging <syslog IPv4 address/IPv6 address/hostname> [trap {<log-level> override class <class> priority <log-level>}]</p> <p>Enables (by setting the syslog IPv4 address/IPv6 address/hostname) sending logging messages, with ability to filter the logging messages according to their classes. The no form of the command stops sending messages to the remote syslog server.</p>	
Syntax Description	syslog IPv4 address/IPv6 address/hostname	<p>syslog IPv4 address/IPv6 address/hostname of the remote syslog server</p> <p>Hostname is limited to 64 characters</p>
	log-level	<ul style="list-style-type: none"> • none—disables the logging locally and remotely • 0 - emerg—system is unusable (emergency) • 1 - alert—alert notification, action must be taken immediately • 2 - crit—critical condition • 3 - err—error condition • 4 - warning—warning condition • 5 - notice—normal, but significant condition • 6 - info—informational condition • 7 - debug—debug level messages
	class	<p>Sets or removes a per-class override on the logging level. All classes which do not have an override set will use the global logging level set with “logging local <log level>”. Classes that do have an override will do as the override specifies. If “none” is specified for the log level, the software will not log anything from this class.</p> <p>Classes available:</p> <ul style="list-style-type: none"> • 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
Default	Remote logging is disabled	
Configuration Mode	config	
History	<p>3.6.2002</p> <p>3.8.1000—Updated command syntax</p>	
Example	switch (config) # logging local info	
Related Commands	<p>show logging</p> <p>logging local override</p> <p>logging <syslog IPv4 address/IPv6 address/hostname> port</p>	
Notes		

8.1.4.4 logging debug-files

	logging debug-files {delete {current oldest} rotation {criteria force max-num} update {<number> current} upload <log-file> <upload URL>} no logging debug-files rotation criteria Configures settings for debug log files. The "logging debug-files rotation criteria" command removes the debug rotation criteria configuration.	
Syntax Description	delete {current oldest}	Deletes certain debug-log files. <ul style="list-style-type: none"> • 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. <ul style="list-style-type: none"> • criteria—sets how the system decides when to rotate debug files <ul style="list-style-type: none"> • 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. <ul style="list-style-type: none"> • current—uploads log file "messages" to a remote host • number—uploads compressed log file "debug.<number>.gz" to a remote host. Range is 1-10.
	upload	Uploads debug log file to a remote host
	log-file	Possible values: 1-7, or current
	upload URL	Supported formats: HTTP, HTTPS, FTP, TFTP, SCP and SFTP (e.g.: scp://username[:password]@hostname/path/filename)
Default	N/A	
Configuration Mode	config	
History	3.3.4150 3.9.0900: Added "no logging debug-files rotation criteria" command	
Example	<pre>switch (config) # logging debug-files delete current</pre>	
Related Commands		
Notes		

8.1.4.5 logging events enable

	logging events {cpu-rate-limiters interfaces protocols} enable no logging events {cpu-rate-limiters interfaces protocols what-just-happened-packets} enable Activate event tracking for a certain group. The no form of the command deactivates event tracking for a certain group.
--	--

Syntax Description	cpu-rate-limiters interfaces protocols what-just-happened-packets	Logical groups with specified set of counters
Default	N/A	
Configuration Mode	config	
History	3.6.6000 3.9.0900: Added note	
Example	switch (config) # logging events interfaces enable	
Related Commands		
Notes	<p>Increase in the enabled events groups will generate a log message of the form:</p> <pre>Jan 8 14:15:24 switch statsd[4404]: [statsd.NOTICE]: (StatsLog) Interface Eth1/9: 398 0598 packets dropped due to Rx invalid tag discards packets Jan 8 14:15:24 switch statsd[4404]: [statsd.NOTICE]: (StatsLog) Interface Eth1/9: 398 0599 packets dropped due to Rx discard packets by vlan filter Jan 8 14:42:44 switch statsd[4404]: [statsd.NOTICE]: (StatsLog) cpu-rate-limiter DISCARD_LAYERS_2_3: 7767087 packets dropped by CPU rate-limiter</pre>	

8.1.4.6 logging events error-threshold

	logging events {interfaces protocols} error-threshold <events> no logging events {interfaces protocols} error-threshold <events> Configures number of events after which the system begins to generate events to the log file. The no form of the command resets this parameter to its default value.	
Syntax Description	interfaces	Sets threshold for interface related events
	protocols	Sets threshold for protocol related events
	events	Number of events after which the system begins to generate events to the log file. Range: 0-4294967295.
Default	cpu-rate-limiters - 1 event interfaces - 10 events protocols - 2 events	
Configuration Mode	config	
History	3.6.6000	
Example	switch (config) # logging events interfaces error-threshold 45	
Related Commands		
Notes		

8.1.4.7 logging events interval

	logging events {interfaces protocols} interval <seconds> no logging events {interfaces protocols} interval <seconds> Configures interval in seconds between each sampling of counters in event type. The no form of the command resets this parameter to its default value.	
Syntax Description	interfaces protocols	Logical groups with specified set of counters Default: interfaces—5 minutes protocols—1 minute
	seconds	Time between sampling. Range is different for each event type: <ul style="list-style-type: none"> • interfaces—10-3600 • protocols—10-3600
Default	N/A	
Configuration Mode	config	
History	3.6.6000	
Example	switch (config) # logging events interfaces interval 120	
Related Commands		
Notes		

8.1.4.8 logging events rate-limit

	logging events [interfaces protocols] rate-limit {short medium long} [count window] no logging events [interfaces protocols] rate-limit [short medium long] [count <number> window <seconds>] Configures the number of allowed events per time window, and that window's duration. The no form of the command resets these parameters to their default values.	
Syntax Description	interfaces protocols	Logical groups with specified set of counters
	rate-limit	Three configurable periods: short, medium, and long
	count	Number of allowed events per time window
	window	Window of time in seconds for the rate limit period
Default	For “interfaces” Short window: event count—5 window duration—1 hour Medium window: event count—50 window duration—1 day Long window: event count—350 window duration—7 days	For “protocols” Short window: event count—10 window duration—1 hour Medium window: event count—100 window duration—1 day Long window: event count—600 window duration—7 days
Configuration Mode	config	
History	3.6.6000	

Example	<code>switch (config) # logging events interfaces interval 120</code>
Related Commands	
Notes	The goal of this command is to restrict the number of events in the log. To achieve this end, it is possible to specify the allowed number (parameter “count”) of messages per period of time (parameter “window”).

8.1.4.9 logging fields

	<code>logging fields seconds {enable fractional-digits <f-digit> whole-digits <w-digit>}</code> <code>no logging fields seconds {enable fractional-digits <f-digit> whole-digits <w-digit>}</code> Specifies whether to include an additional field in each log message that shows the number of seconds since the Epoch or not. The no form of the command disallows including an additional field in each log message that shows the number of seconds since the Epoch.	
Syntax Description	enable	Specifies whether to include an additional field in each log message that shows the number of seconds since the Epoch or not.
	f-digit	The fractional-digits parameter controls the number of digits to the right of the decimal point. Truncation is done from the right. Possible values are: 1, 2, 3, or 6.
	w-digit	The whole-digits parameter controls the number of digits to the left of the decimal point. Truncation is done from the left. Except for the year, all of these digits are redundant with syslog's own date and time. Possible values: 1, 6, or all.
Default	Disabled	
Configuration Mode	config	
History	3.1.0000	
Example	<code>switch (config) # logging fields seconds enable</code> <code>switch (config) # logging fields seconds whole-digits 1</code>	
Related Commands	show logging	
Notes	This is independent of the standard syslog date and time at the beginning of each message in the format of “July 15 18:00:00”. Aside from indicating the year at full precision, its main purpose is to provide subsecond precision.	

8.1.4.10 logging files delete

	<code>logging files delete {current oldest [<number of files>]}</code> Deletes the current or oldest log files.	
Syntax Description	current	Deletes current log file
	oldest	Deletes oldest log file
	number of files	Sets the number of files to be deleted
Default	CLI commands and audit message are set to notice logging level	
Configuration Mode	config	
History	3.1.0000	

Example	<code>switch (config) # logging files delete current</code>
Related Commands	<code>show logging</code> <code>show log files</code>
Notes	

8.1.4.11 logging files rotation

	<code>logging files rotation {criteria {frequency <freq> size <size-mb> size-pct <size-percentage>} force max-number <number-of-files>}</code> <code>no logging files rotation criteria</code> Sets the rotation criteria of the logging files. The no form of the command removes the rotation criteria configuration.	
Syntax Description	freq	Sets rotation criteria according to time. Possible options are: <ul style="list-style-type: none"> • Daily • Weekly • Monthly
	size-mb	Sets rotation criteria according to size in megabytes Range: 1-9999 Default: 20MB
	size-percentage	Sets rotation criteria according to size in percentage of the partition where the logging files are kept in. The percentage given is truncated to three decimal points (thousandths of a percent).
	force	Forces an immediate rotation of the log files. This does not affect the schedule of auto-rotation if it was done based on time: the next automatic rotation will still occur at the same time for which it was previously scheduled. Naturally, if the auto-rotation was based on size, this will delay it somewhat as it reduces the size of the active log file to zero.
	number-of-files	The number of log files will be kept. If the number of log files ever exceeds this number (either at rotation time, or when this setting is lowered), the system will delete as many files as necessary to bring it down to this number, starting with the oldest.
Default	10 files are kept by default with rotation criteria of 5% of the log partition size	
Configuration Mode	config	
History	3.1.0000 3.9.0900: <ul style="list-style-type: none"> • Added the command "no logging files rotation criteria" • Changed default value size from 19.07 MB to 20 MB 	
Example	<code>switch (config) # logging files rotation criteria size-pct 6</code>	
Related Commands	<code>show logging</code> <code>show log files</code>	
Notes		

8.1.4.12 logging files upload

	<code>logging files upload {current <file-number>} <url></code> Uploads a log file to a remote host.
--	---

Syntax Description	current	The current log file. The current log file will have the name “messages” if you do not specify a new name for it in the upload URL.
	file-number	An archived log file. The archived log file will have the name “messages<n>.gz” (while “n” is the file number) if you do not specify a new name for it in the upload URL. The file will be compressed with gzip.
	url	Uploads URL path. Supported formats: FTP, TFTP, SCP, and SFTP. For example: scp://username[:password]@hostname/path/filename.
Default	10 files are kept by default with rotation criteria of 5% of the log partition size	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # logging files upload 1 scp://admin@scpserver	
Related Commands	show logging show log files	
Notes		

8.1.4.13 logging filter include

	logging <IP address>\hostname> filter include <regex> Sends only log messages that match the input regex to a remote host specified by its IP or hostname.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config	
History	3.8.2000	
Role	admin	
Example	switch (config) # logging 1.1.1.1 filter include ERROR	
Related Commands	loggin no logging	
Notes	This command is configurable. If “configuration write” is executed, the remote host will still receive filtered messages after reload.	

8.1.4.14 logging filter exclude

	logging <IP address>\hostname> filter exclude <regex> Sends only log messages that do not match the input regex to a remote host specified by its IP or hostname.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config	
History	3.8.2000	
Role	admin	

Example	<code>switch (config) # logging 1.1.1.1 filter exclude ERROR</code>
Related Commands	<code>loggin</code> <code>no logging</code>
Notes	This command is configurable. If “configuration write” is executed, the remote host will still receive filtered messages after reload.

8.1.4.15 no logging filter

	<code>no logging <IP address\hostname> filter</code> Sends unfiltered log messages to the configured remote host.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.8.2000
Role	admin
Example	<code>switch (config) # no logging 1.1.1.1 filter</code>
Related Commands	<code>loggin</code> <code>no logging</code>
Notes	This command is configurable. If “configuration write” is executed, the remote host will still receive filtered messages after reload.

8.1.4.16 logging format

	<code>logging format {standard welf [fw-name <hostname>]}</code> <code>no logging format {standard welf [fw-name <hostname>]}</code> Sets the format of the logging messages. The no form of the command resets the format to its default.	
Syntax Description	standard	Standard format
	welf	WebTrends Enhanced Log file (WELF) format
	hostname	Specifies the firewall hostname that should be associated with each message logged in WELF format. If no firewall name is set, the hostname is used by default. Hostname is limited to 64 characters.
Default	standard	
Configuration Mode	config	
History	3.1.0000	
Example	<code>switch (config) # logging format standard</code>	
Related Commands	<code>show logging</code>	
Notes		

8.1.4.17 logging level

	logging level {cli commands <log-level> audit mgmt <log-level>} Sets the severity level at which CLI commands or the management audit message that the user executes are logged. This includes auditing of both configuration changes and actions.	
Syntax Description	cli commands	Sets the severity level at which CLI commands which the user executes are logged
	audit mgmt	Sets the severity level at which all network management audit messages are logged
	log-level	<ul style="list-style-type: none"> • none—disables the logging locally and remotely • 0 - emerg—system is unusable (emergency) • 1 - alert—alert notification, action must be taken immediately • 2 - crit—critical condition • 3 - err—error condition • 4 - warning—warning condition • 5 - notice—normal, but significant condition • 6 - info—informational condition • 7 - debug—debug level messages
Default	CLI commands and audit message are set to notice logging level	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # logging level cli commands info	
Related Commands	show logging	
Notes		

8.1.4.18 logging local override

	logging local override [class <class> priority <log-level>] no logging local override [class <class> priority <log-level>] Enables class-specific overrides to the local log level. The no form of the command disables all class-specific overrides to the local log level without deleting them from the configuration, but disables them so that the logging level for all classes is determined solely by the global setting.	
Syntax Description	override	Enables class-specific overrides to the local log level.
	class	Sets or removes a per-class override on the logging level. All classes which do not have an override set will use the global logging level set with “logging local <log level>”. Classes that do have an override will do as the override specifies. If “none” is specified for the log level, the software will not log anything from this class. Classes available: <ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • none—disables the logging locally and remotely • 0 - emerg—system is unusable (emergency) • 1 - alert—alert notification, action must be taken immediately • 2 - crit—critical condition • 3 - err—error condition • 4 - warning—warning condition • 5 - notice—normal, but significant condition • 6 - info—informational condition • 7 - debug—debug level messages
Default	Override is disabled	
Configuration Mode	config	
History	3.1.0000 3.3.4150: Added debug-module class and changed iss-modules to protocol-stack	
Example	switch (config) # logging local override class mgmt-front priority warning	
Related Commands	show logging logging local	
Notes		

8.1.4.19 logging monitor

	<p>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.</p>	
Syntax Description	facility	<ul style="list-style-type: none"> • mgmt-front • mgmt-back • mgmt-core • events • sx-sdk • mlnx-daemons • iss-modules
	priority-level	<ul style="list-style-type: none"> • none • emerg • alert • crit • err • warning • notice • info • debug
Default	no logging monitor	
Configuration Mode	config	
History	3.3.4000	
Example	switch (config) # logging monitor events notice	
Related Commands		
Notes		

8.1.4.20 logging protocol

	logging <IP address\hostname> protocol [tcp udp] no logging <IP address\hostname> protocol Sends log messages to specified host with the chosen protocol (TCP or UDP). The no form of the command sets the protocol for sending log messages to a remote host to the default (UDP).	
Syntax Description	tcp	Sets protocol to TCP
	udp	Sets protocol to UDP
Default	UDP	
Configuration Mode	Configure terminal	
History	3.8.2100	
Role	Admin	
Example	<pre>switch (config) # logging 1.1.1.1 protocol tcp switch (config) # no logging 1.1.1.1 protocol</pre>	
Related Commands		
Notes	This command is configurable, so if “configuration write” is executed then after reboot the remote host will still receive messages with the configured protocol.	

8.1.4.21 logging receive

	logging receive no logging receive Enables receiving logging messages from a remote host. The no form of the command disables the option of receiving logging messages from a remote host.	
Syntax Description	N/A	
Default	Receiving logging is disabled	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # logging receive</pre>	
Related Commands	show logging logging local logging local override	
Notes	<ul style="list-style-type: none"> • This does not log to the console TTY port • In-band management should be enabled in order to open a channel from the host to the CPU • If enabled, only log messages matching or exceeding the minimum severity specified with the “logging local” command will be logged, regardless of what is sent from the remote host 	

8.1.4.22 logging mac masking

	logging mac masking no logging mac masking Enables MAC address masking in logs. The no form of the command disables MAC address masking.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config
History	3.9.0900
Example	switch (config) # logging mac masking
Related Commands	show logging
Notes	If enabled, the first 2 bytes of MAC address output log will be masked. For example, 00:12:34:56:78:9a will be displayed as **:*.34:56:78:9a.

8.1.4.23 show log

	show log [continuous files [<file-number>]] [[not] matching <reg-exp>] Displays the log file with optional filter criteria.	
Syntax Description	continues	Displays the last few lines of the current log file and then continues to display new lines as they come in until the user hits Ctrl+C, similar to LINUX “tail” utility
	files	Displays the list of log files
	<file-number>	Displays an archived log file, where the number may range from 1 up to the number of archived log files available
	[not] matching <reg-exp>	The file is piped through a LINUX “grep” utility to only include lines either matching, or not matching, the provided regular expression
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000 3.3.4402: Updated example and added note	
Example	<pre>switch (config) # show log matching "Executing Action" Jul 31 16:11:23 M2100-aj cli[26502]: [cli.NOTICE]: user : Executing command: enable Jul 31 16:11:24 M2100-aj cli[26507]: [cli.NOTICE]: user : Executing command: enable Jul 31 16:11:29 M2100-aj cli[26514]: [cli.NOTICE]: user : Executing command: enable Jul 31 16:11:29 M2100-aj cli[26514]: [cli.NOTICE]: user : Executing command: show license Jul 31 16:11:41 M2100-aj cli[26548]: [cli.NOTICE]: user : Executing command: enable Jul 31 16:11:42 M2100-aj cli[26553]: [cli.NOTICE]: user : Executing command: enable Jul 31 16:11:42 M2100-aj cli[26553]: [cli.NOTICE]: user : Executing command: conf termina</pre>	
Related Commands	logging fields logging files rotation logging level logging local logging receive show logging	

Notes	<ul style="list-style-type: none"> When using a regular expression containing (OR), the expression should be surrounded by quotes (“<expression>”), otherwise it is parsed as filter (PIPE) command The command’s output has many of the options as the Linux “less” command. These options allow navigating the log file and perform searches. To see help for different option press “h” after running the “show log” command.
-------	--

8.1.4.24 show logging

	show logging Displays the logging configurations.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.8.2000: Updated example 3.9.0900: Updated example
Role	Admin
Example	<pre> switch (config) # show logging Local logging level : notice Override for class debug-module : notice Default remote logging level : notice Allow receiving of messages from remote hosts: no Number of archived log files to keep : 10 Log rotation size threshold : 19.07 megabytes Log rotation (debug) size threshold : 19.07 megabytes Log format : standard Subsecond timestamp field : disabled MAC address masking : enabled Levels at which messages are logged: CLI commands : notice Audit messages: notice Remote syslog servers: 1.1.1.1: log level : notice Remote port : 514 Filter [include] regex: err 1.2.2.3: log level : notice Remote port: 33 </pre>
Related Commands	logging fields logging files rotation logging level logging local logging receive logging <syslog IPv4 address/IPv6 address/hostname>
Notes	

8.1.4.25 show logging events

	show logging events [interfaces protocols] Displays configuration per selected event group or all.	
Syntax Description	interfaces protocols	Logical groups with specified set of counters
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.6000	
Example	<pre>switch (config) # show logging events interfaces: Admin mode : no Interval : 5 minutes Error threshold: 10 Rate-limit short window: Event count : 5 Window duration: 1 hour Rate-limit medium window: Event count : 50 Window duration: 1 day Rate-limit long window: Event count : 350 Window duration: 7 days protocols: Admin mode : no Interval : 1 minute Error threshold: 2 Rate-limit short window: Event count : 10 Window duration: 1 hour Rate-limit medium window: Event count : 100 Window duration: 1 day Rate-limit long window: Event count : 600 Window duration: 7 days</pre>	
Related Commands	logging event enable logging event error-threshold logging event interval logging event rate-limit	
Notes		

8.1.4.26 show logging events source-counters

	show logging events [interfaces protocols] source-counters Displays set of counters for sampling.	
Syntax Description	interfaces protocols	Logical groups with specified set of counters
Default	N/A	

Configuration Mode	Any command mode
History	3.6.6000
Example	<pre>switch (config) # show logging events interfaces source-counters interfaces: Counters: Rx discard packets, Rx error packets, Rx fcs errors, Rx undersize packets, Rx oversize packets, Rx unknown control opcode, Rx symbol errors, Rx discard packets by Storm Control, Tx discard packets, Tx error packets, Tx hog discard packets</pre>
Related Commands	<pre>logging event enable logging event error-threshold logging event interval logging event rate-limit</pre>
Notes	

8.1.4.27 show logging port

	<pre>show logging port Displays the port logging configurations.</pre>
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	<pre>3.1.0000 3.8.1000: Updated example</pre>
Example	<pre>switch (config) # show logging Local logging level: notice Override for class debug-module: notice Default remote logging level: notice Remote syslog receiver: 1.2.3.4 (log level: notice) Remote port: 514</pre>
Related Commands	logging port
Notes	

8.2 Link Diagnostic Per Port



When debugging a system, it is important to be able to quickly identify the root of a problem. The Diagnostic commands enables an insight into the physical layer components where the user is able to see information such as a cable status (plugged/unplugged) or if Auto-Negotiation has failed.

List of possible output messages:

- 0—No issue observed
- 1—Port is close by command (see PAOS)
- 2—AN no partner detected
- 3—AN ack not received
- 4—AN next-page exchange failed
- 5—KR frame lock not acquired
- 6—KR link inhibit timeout

7–KR Link partner didn't set receiver ready
 8–KR tuning didn't completed
 9–PCS didn't acquire block lock
 10–PCS didn't acquire AM lock (NO FEC)
 11–PCS didn't get align_status
 12–FC FEC is not locked
 13–RS FEC is not locked
 14–Remote fault received
 15–Bad Signal integrity
 16–Compliance code mismatch (protocol mismatch between cable and port)
 17–Large number of physical errors (high BER)
 18–Port is disabled by Ekey
 19–Phase EO failure
 20–Stamping of non-NVIDIA Cables/Modules
 21–Down by PortInfo MAD
 22–Disabled by Verification
 23–Calibration failure
 24–EDR speed is not allowed due to cable stamping: EDR stamping
 25–FDR10 speed is not allowed due to cable stamping: FDR10 stamping
 26–Port is closed due to cable stamping: Ethernet_compliance_code_zero
 27–Port is closed due to cable stamping: 56GE stamping
 28–Port is closed due to cable stamping: non-NVIDIA QSFP28
 29–Port is closed due to cable stamping: non-NVIDIA SFP28
 30–Port is closed, no backplane enabled speed over backplane channel
 31–Port is closed, no passive protocol enabled over passive copper channel
 32–Port is closed, no active protocol enabled over active channel
 33–Port width is does not match the port speed enabled
 34–Local speed degradation
 35–Remote speed degradation
 36–No Partner detected during force mode.
 37–Partial link indication during force mode.
 38–AN Failure–FEC mismatch during override
 39–AN Failure–No HCD
 40–VPI protocol don't match
 41–Port is closed, module can't be set to the enabled rate
 42–Bad SI, cable is configured to non optimal rate
 1023–Info not available
 MNG FW issues (1024–2047):
 1024–Cable is unplugged/powering off
 1025–Long Range for non MLNX cable/module .
 1026–Bus stuck (I2C Data or clock shorted)
 1027–bad/unsupported EEPROM
 1028–part number list
 1029–unsupported cable.
 1030–module temperature shutdown
 1031–Shorted cable
 1032–Power Budget Exceeded
 1033–Management force down the port
 1034–Module is disabled by command

8.2.1 Link Diagnostic Commands

8.2.1.1 show interfaces ib link-diagnostics

	<p>show interfaces ib [device/port] link-diagnostics</p> <p>Displays a specific InfiniBand module/port or all InfiniBand ports.</p>
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.4000
Example	<pre>switch (config) # show interfaces ib link-diagnostics ----- Interface Code Status ----- IB1/1 0 The port is Active. IB1/2 0 The port is Active. IB1/3 1024 Cable unplugged IB1/4 1024 Cable unplugged IB1/5 1024 Cable unplugged IB1/6 1024 Cable unplugged IB1/7 1024 Cable unplugged IB1/8 1024 Cable unplugged IB1/9 1024 Cable unplugged IB1/10 1024 Cable unplugged IB1/11 1024 Cable unplugged IB1/12 1024 Cable unplugged IB1/13 1024 Cable unplugged IB1/14 1024 Cable unplugged IB1/15 1024 Cable unplugged IB1/16 1024 Cable unplugged IB1/17 1024 Cable unplugged IB1/18 1024 Cable unplugged IB1/19 1024 Cable unplugged IB1/20 1024 Cable unplugged IB1/21 1024 Cable unplugged IB1/22 1024 Cable unplugged IB1/23 1024 Cable unplugged IB1/24 1024 Cable unplugged IB1/25 1024 Cable unplugged IB1/26 1024 Cable unplugged IB1/27 1024 Cable unplugged IB1/28 1024 Cable unplugged IB1/29 1024 Cable unplugged IB1/30 1024 Cable unplugged IB1/31 1024 Cable unplugged IB1/32 1024 Cable unplugged IB1/33 1024 Cable unplugged IB1/34 1024 Cable unplugged IB1/35 1 The port is closed by command. IB1/36 2 Auto-Negotiation failure..</pre>
Related Commands	
Notes	

8.2.1.2 show interfaces ib internal leaf link-diagnostics

	<p>show interfaces ib internal leaf <module/port> link-diagnostics</p> <p>Displays a specific InfiniBand internal leaf module/port.</p>
Syntax Description	N/A

Default	N/A
Configuration Mode	config
History	3.6.4000
Example	<pre>switch (config) # show interfaces ib internal leaf 1 link-diagnostics ----- Interface Code Status ----- IB1/1/19 0 No issue was observed IB1/1/20 0 No issue was observed IB1/1/21 0 No issue was observed IB1/1/22 0 No issue was observed IB1/1/23 0 No issue was observed IB1/1/24 0 No issue was observed IB1/1/25 0 No issue was observed IB1/1/26 0 No issue was observed IB1/1/27 0 No issue was observed IB1/1/28 0 No issue was observed IB1/1/29 0 No issue was observed IB1/1/30 0 No issue was observed</pre>
Related Commands	
Notes	

8.2.1.3 show interfaces ib internal spine link-diagnostics

	<pre>show interfaces ib internal spine <module/port> link-diagnostics Displays a specific InfiniBand internal spine module/port.</pre>
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.4000
Example	<pre>switch (config) # show interfaces ib internal spine 3/1/1 link-diagnostics ----- Interface Code Status ----- IB3/1/1 0 No issue was observed</pre>
Related Commands	
Notes	

8.3 Event Notifications



The OS features a variety of supported events. Events are printed in the system log file and can, optionally, be sent to the system administrator via email, SNMP trap or directly prompted to the terminal.

8.3.1 Supported Event Notifications and MIB Mapping

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

Event Name	Event Description	MIB OID	Comments
asic-chip-down	ASIC (chip) down	Mellanox-EFM-MIB: asicChipDown	Not supported
cpu-util-high	CPU utilization has risen too high	Mellanox-EFM-MIB: cpuUtilHigh	N/A
disk-space-low	File system free space has fallen too low	Mellanox-EFM-MIB: diskSpaceLow	N/A
health-module-status	Health module status changed	Mellanox-EFM-MIB: systemHealthStatus	N/A
insufficient-fans	Insufficient amount of fans in system	Mellanox-EFM-MIB: insufficientFans	N/A
insufficient-fans-recover	Insufficient amount of fans in system recovered	Mellanox-EFM-MIB: insufficientFansRecover	N/A
insufficient-power	Insufficient power supply	Mellanox-EFM-MIB: insufficientPower	N/A
interface-down	An interface's link state has changed to DOWN	RFC1213: linkdown (SNMPv1)	Supported for InfiniBand interfaces for 1U and blade systems
interface-up	An interface's link state has changed to UP	RFC1213: linkup (SNMPv1)	Supported for InfiniBand interfaces for 1U and blade systems
internal-bus-error	Internal bus (I2C) error	Mellanox-EFM-MIB: internalBusError	N/A
internal-link-speed-mismatch	There is a mismatch in the speeds of the internal links between spine and leaf modules	Mellanox-EFM-MIB: internalSpeedMismatch	Supported only for modular switches
liveness-failure	A process in the system is detected as hung	Not implemented	N/A
low-power	Low power supply	Mellanox-EFM-MIB: lowPower	N/A
low-power-recover	Low power supply recover	Mellanox-EFM-MIB: lowPowerRecover	N/A
paging-high	Paging activity has risen too high	N/A	Not supported
power-redundancy-mismatch	Power redundancy mismatch	Mellanox-EFM-MIB: powerRedundancyMismatch	Supported only for modular switches
process-crash	A process in the system has crashed	Mellanox-EFM-MIB: procCrash	N/A
process-exit	A process in the system unexpectedly exited	Mellanox-EFM-MIB: procUnexpectedExit	N/A
send-test	Send a test notification	testTrap	Run the CLI command "snmp-server notify send-test"
snmp-authtrap	An SNMPv3 request has failed authentication	Not implemented	N/A

Event Name	Event Description	MIB OID	Comments
temperature-too-high	Temperature is too high	Mellanox-EFM-MIB:asicOverTemp	N/A
unexpected-shutdown	Unexpected system shutdown	Mellanox-EFM-MIB:unexpectedShutdown	N/A
cli-line-executed			
disk-io-high			
entity-state-change			
expected-shutdown			
memusage-high			
netusage-high			
sm-restart			
sm-start			
sm-stop			
unexpected-cluster-join			
unexpected-cluster-leave			
unexpected-cluster-size			
user-login			
user-logout			

8.3.2 SNMP Trap Notification

To set SNMP notification see [Configuring SNMP Notifications \(Traps or Informs\)](#) section.

8.3.3 Terminal Notifications

To print events to the terminal, set the events you wish to print to the terminal. Run:

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

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

```
switch (config) #
Wed Jul 10 11:30:42 2022: Interface 1/17 changed state to DOWN
Wed Jul 10 11:30:43 2022: Interface 1/18 changed state to DOWN
switch (config) #
```

8.3.4 Email Notifications

To configure the OS to send you emails for all configured events and failures:

1. Set your mailhub to the IP address to be your mail client's server - for example, Microsoft Outlook exchange server.

```
switch (config) # email mailhub <IP address>
```

2. Add your email address for notifications. Run:

```
switch (config) # email notify recipient <email address>
```

3. Configure the system to send notifications for a specific event. Run:

```
switch (config) # email notify event <event name>
```

4. Show the list of events for which an email is sent. Run:

```
switch (config) # show email events
Failure events for which emails will be sent:
  process-crash: A process in the system has crashed
  unexpected-shutdown: Unexpected system shutdown

Informational events for which emails will be sent:
  asic-chip-down: ASIC (Chip) Down
  cpu-util-high: CPU utilization has risen too high
  cpu-util-ok: CPU utilization has fallen back to normal levels
  disk-io-high: Disk I/O per second has risen too high
  disk-io-ok: Disk I/O per second has fallen back to acceptable levels
  disk-space-low: Filesystem free space has fallen too low
...
```

5. Have the system send you a test email. Run:

```
switch (config) # email send-test

The last command should generate the following email:
-----Original Message-----
From: Admin User [mailto:do-not-reply@switch.]
Sent: Sunday, May 01, 2011 11:17 AM
To: <name>
Subject: System event on switch: Test email for event notification

==== System information:
Hostname: switch

Version: MLNX-OS 3.11.1954-007 #1-dev 2023-10-18 15:21:05
Date: 2023/10/19 16:34:04
Uptime: 1h 45m 8.730s

This is a test email.

==== Done.
```

8.3.5 Command Event Notifications

8.3.5.1 email autosupport enable

	email autosupport enable no email autosupport enable Sends automatic support notifications via email. The no form of the command stops sending automatic support notifications via email.
Syntax Description	N/A

Default	N/A
Configuration Mode	config
History	3.2.3000
Example	<code>switch (config) # email autosupport enable</code>
Related Commands	
Notes	

8.3.5.2 email autosupport event

	<p>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.</p>	
Syntax Description	event	<ul style="list-style-type: none"> • process-crash - a process has crashed • process-exit - a process unexpectedly exited • liveness-failure - a process is detected as hung • cpu-util-high - CPU utilization has risen too high • cpu-util-ok - CPU utilization has fallen back to normal levels • paging-high - paging activity has risen too high • paging-ok - paging activity has fallen back to normal levels • disk-space-low - filesystem free space has fallen too low • disk-space-ok - filesystem free space is back in the normal range • memusage-high - memory usage has risen too high • memusage-ok - memory usage has fallen back to acceptable levels • netusage-high - network utilization has risen too high • netusage-ok - network utilization has fallen back to acceptable levels • disk-io-high - disk I/O per second has risen too high • disk-io-ok - disk I/O per second has fallen back to acceptable levels • unexpected-cluster-join - node has unexpectedly joined the cluster • unexpected-cluster-leave - node has unexpectedly left the cluster • unexpected-cluster-size - the number of nodes in the cluster is unexpected • unexpected-shutdown - unexpected system shutdown • interface-up - an interface's link state has changed to up • interface-down - an interface's link state has changed to down • user-login - a user has logged into the system • user-logout - a user has logged out of the system • health-module-status - health module status • temperature-too-high - temperature has risen too high • low-power - low power supply • low-power-recover - low power supply recover • insufficient-power - insufficient power supply • power-redundancy-mismatch - power redundancy mismatch • insufficient-fans - insufficient amount of fans in system • insufficient-fans-recover - insufficient amount of fans in system recovered • asic-chip-down - ASIC (chip) down • internal-bus-error - internal bus (I2C) error • internal-link-speed-mismatch - internal links speed mismatch

Default	N/A
Configuration Mode	config
History	3.2.3000
Example	switch (config) # email autosupport event process-crash
Related Commands	
Notes	

8.3.5.3 email autosupport ssl mode

	email autosupport ssl mode {none tls tls-none} no email autosupport ssl mode Configures type of security to use for auto-support email. The no form of the command resets auto-support email security mode to its default.	
Syntax Description	none	Does not use TLS to secure auto-support email.
	tls	Uses TLS over the default server port to secure auto-support email and does not send an email if TLS fails.
	tls-none	Attempts TLS over the default server port to secure auto-support email, and falls back on plaintext if this fails.
Default	tls-none	
Configuration Mode	config	
History	3.2.3000	
Example	switch (config) # email autosupport ssl mode tls	
Related Commands		
Notes		

8.3.5.4 email autosupport ssl cert-verify

	email autosupport ssl cert-verify no email autosupport ssl cert-verify Verifies server certificates. The no form of the command does not verify server certificates.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Example	switch (config) # email autosupport ssl cert-verify	
Related Commands		
Notes		

8.3.5.5 email autosupport ssl ca-list

	email autosupport ssl ca-list {<ca-list-name> default_ca_list none} no email autosupport ssl ca-list Configures supplemental CA certificates for verification of server certificates. The no form of the command removes supplemental CA certificate list.	
Syntax Description	default_ca_list	Default supplemental CA certificate list
	none	No supplemental list (uses built-in list only)
Default	default_ca_list	
Configuration Mode	config	
History	3.2.3000	
Example	switch (config) # email autosupport ssl ca-list default_ca_list	
Related Commands		
Notes		

8.3.5.6 email dead-letter

	email dead-letter {cleanup max-age <duration> enable} no email dead-letter Configures settings for saving undeliverable emails. The no form of the command disables sending of emails to vendor auto-support upon certain failures.	
Syntax Description	duration	Example: "5d4h3m2s" for 5 days, 4 hours, 3 minutes, 2 seconds
	enable	Saves dead-letter files for undeliverable emails
Default	Save dead letter is enabled The default duration is 14 days	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # email dead-letter enable	
Related Commands	show email	
Notes		

8.3.5.7 email domain

	email domain <hostname-or-ip-address> no email domain Sets the domain name from which the emails appear to come (provided that the return address is not already fully-qualified). This is used in conjunction with the system hostname to form the full name of the host from which the email appears to come. The no form of the command clears email domain override.	
Syntax Description	hostname-or-ip-address	Hostname or IP address of email domain
Default	No email domain	

Configuration Mode	config
History	3.1.0000
Example	switch (config) # email domain my_domain
Related Commands	show emails
Notes	

8.3.5.8 email mailhub

	email mailhub <hostname-or-ip-address> no email mailhub Sets the mail relay to be used to send notification emails. The no form of the command clears the mail relay to be used to send notification emails.	
Syntax Description	hostname-or-ip-address	Hostname or IP address
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # email mailhub 10.0.8.11	
Related Commands	show email [events]	
Notes		

8.3.5.9 email autosupport mailhub

	email autosupport mailhub <hostname-or-ip-address> no email autosupport mailhub Sets the mail relay to be used for sending autosupport notification emails. The no form of the command clears the mail relay to be used for sending autosupport notification emails.	
Syntax Description	<hostname-or-ip-address>	The mail hub hostname or IP address
Default	N/A	
Configuration Mode	config	
History	3.7.1000	
Example	switch (config) # email autosupport mailhub 10.10.10.1	
Related Commands	show email	
Notes		

8.3.5.10 email autosupport recipient

	email autosupport recipient <email-addr> no email autosupport recipient Sets the recipient for autosupport emails. The no form of the command clears the configured autosupport recipient.	
Syntax Description	email-addr	The autosupport recipient email address
Default	N/A	
Configuration Mode	config	
History	3.7.1000	
Example	switch (config) # email autosupport recipient user@example.com	
Related Commands	show email	
Notes		

8.3.5.11 email mailhub-port

	email mailhub-port <port number> no email mailhub-port Sets the mail relay port to be used to send notification emails. The no form of the command resets the port to its default.	
Syntax Description	hostname-or-ip-address	Port number
Default	25	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # email mailhub-port 125	
Related Commands	show email	
Notes		

8.3.5.12 email notify event

	email notify event <event> no email notify event <event> Enables sending email notifications for the specified event type. The no form of the command disables sending email notifications for the specified event type.	
--	---	--

Syntax Description	event	<p>Available event names:</p> <ul style="list-style-type: none"> • process-crash - a process has crashed • process-exit - a process unexpectedly exited • liveness-failure - a process is detected as hung • cpu-util-high - CPU utilization has risen too high • cpu-util-ok - CPU utilization has fallen back to normal levels • paging-high - paging activity has risen too high • paging-ok - paging activity has fallen back to normal levels • disk-space-low - filesystem free space has fallen too low • disk-space-ok - filesystem free space is back in the normal range • memusage-high - memory usage has risen too high • memusage-ok - memory usage has fallen back to acceptable levels • netusage-high - network utilization has risen too high • netusage-ok - network utilization has fallen back to acceptable levels • disk-io-high - disk I/O per second has risen too high • disk-io-ok - disk I/O per second has fallen back to acceptable levels • unexpected-cluster-join - node has unexpectedly joined the cluster • unexpected-cluster-leave - node has unexpectedly left the cluster • unexpected-cluster-size - the number of nodes in the cluster is unexpected • unexpected-shutdown - unexpected system shutdown • interface-up - an interface's link state has changed to up • interface-down - an interface's link state has changed to down • user-login - a user has logged into the system • user-logout - a user has logged out of the system • health-module-status - health module status • temperature-too-high - temperature has risen too high • low-power - low power supply • low-power-recover - low power supply recover • insufficient-power - insufficient power supply • power-redundancy-mismatch - power redundancy mismatch • insufficient-fans - insufficient amount of fans in system • insufficient-fans-recover - insufficient amount of fans in system recovered • asic-chip-down - ASIC (chip) down • internal-bus-error - internal bus (I2C) error • internal-link-speed-mismatch - internal links speed mismatch
Default	No events are enabled	
Configuration Mode	config	
History	3.1.0000	
Example	<code>switch (config) # email notify event process-crash</code>	
Related Commands	<p>email autosupport event show email show email events</p>	
Notes	This does not affect auto-support emails. Auto-support can be disabled overall, but if it is enabled, all auto-support events are sent as emails.	

8.3.5.13 email notify recipient

	<p>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.</p>	
Syntax Description	email-addr	Email address of intended recipient.
	class	Specifies which types of events are sent to this recipient.
	info	Sends informational events to this recipient.
	failure	Sends failure events to this recipient.
	detail	Sends detailed event emails to this recipient.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # email notify recipient user2@autosupport.mydomain.com	
Related Commands	show email	
Notes		

8.3.5.14 email return-addr

	<p>email return-addr <username> no email domain Sets the username or fully-qualified return address from which email notifications are sent.</p> <ul style="list-style-type: none"> • If the string provided contains an “@” character, it is considered to be fully-qualified and used as-is. • Otherwise, it is considered to be just the username, and we append “@<hostname>.<domain>”. The default is “do-not-reply”, but this can be changed to “admin” or whatnot in case something along the line does not like fictitious addresses. <p>The no form of the command resets this attribute to its default.</p>	
Syntax Description	username	Username
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # email return-addr user1	
Related Commands	show email	
Notes		

8.3.5.15 email return-host

	<p>email return-host no email return-host</p> <p>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.</p>
Syntax Description	N/A
Default	No return host
Configuration Mode	config
History	3.1.0000
Example	<code>switch (config) # no email return-host</code>
Related Commands	show email
Notes	This only takes effect if the return address does not contain an “@” character

8.3.5.16 email send-test

	<p>email send-test</p> <p>Sends test-email to all configured event and failure recipients.</p>
Syntax Description	N/A
Default	No return host
Configuration Mode	config
History	3.1.0000
Example	<code>switch (config) # email send-test</code>
Related Commands	show email [events]
Notes	

8.3.5.17 email ssl mode

	<p>email ssl mode {none tls tls-none}</p> <p>no email ssl mode</p> <p>Sets the security mode(s) to try for sending email. The no form of the command resets the email SSL mode to its default.</p>	
Syntax Description	none	No security mode, operates in plaintext
	tls	Attempts to use TLS on the regular mailhub port, with STARTTLS. If this fails, it gives up.
	tls-none	Attempts to use TLS on the regular mailhub port, with STARTTLS. If this fails, it falls back on plaintext.
Default	default-cert	
Configuration Mode	config	
History	3.2.3000	
Example	<code>switch (config) # email ssl mode tls-none</code>	

Related Commands	show email
Notes	

8.3.5.18 email ssl cert-verify

	<pre>email ssl cert-verify no email ssl cert-verify</pre> <p>Enables verification of SSL/TLS server certificates for email. The no form of the command disables verification of SSL/TLS server certificates for email.</p>
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.2.3000
Example	<pre>switch (config) # email ssl cert-verify</pre>
Related Commands	show email
Notes	This command has no impact unless TLS is used.

8.3.5.19 email ssl ca-list

	<pre>email ssl ca-list {<ca-list-name> default-ca-list none} no email ssl ca-list</pre> <p>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.</p>	
Syntax Description	ca-list-name	Specifies CA list name
	default-ca-list	Uses default supplemental CA certificate list
	none	Uses no list of supplemental certificates
Default	default-ca-list	
Configuration Mode	config	
History	3.2.3000	
Example	<pre>switch (config) # email ssl ca-list none</pre>	
Related Commands	show email	
Notes	This command has no impact unless TLS is used, and certificate verification is enabled.	

8.3.5.20 show email

	<pre>show email</pre> <p>Displays email configuration or events for which email should be sent upon.</p>
Syntax Description	N/A

Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre> switch (config) # show email Mail hub: 10.0.8.70 Mail hub port: 25 Domain override: Return address: do-not-reply Include hostname in return address: yes Current reply address: do-not-reply@<hostname> Security mode: tls-none Verify server cert: yes Supplemental CA list: default-ca-list Dead letter settings: Save dead.letter files: yes Dead letter max age: 14 days Email notification recipients: No recipients configured. Autosupport emails Enabled: no Recipient: Mail hub: Security mode: tls-none Verify server cert: yes Supplemental CA list: default-ca-list </pre>
Related Commands	
Notes	

8.3.5.21 show email events

	show email events Displays list of events for which notification emails are sent.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000

<p>Example</p>	<pre>switch (config) # show email events Failure events for which emails will be sent: expected-shutdown: Expected system shutdown process-crash: A process in the system has crashed unexpected-shutdown: Unexpected system shutdown Informational events for which emails will be sent: asic-chip-down: ASIC (Chip) Down cpu-util-high: CPU utilization has risen too high cpu-util-ok: CPU utilization has fallen back to normal levels disk-io-high: Disk I/O per second has risen too high disk-io-ok: Disk I/O per second has fallen back to acceptable levels disk-space-low: Filesystem free space has fallen too low disk-space-ok: Filesystem free space is back in the normal range health-module-status: Health module Status insufficient-fans: Insufficient amount of fans in system insufficient-fans-recover: Insufficient amount of fans in system recovered insufficient-power: Insufficient power supply internal-bus-error: Internal bus (I2C) Error internal-link-speed-mismatch: Internal links speed mismatch liveness-failure: A process in the system was detected as hung low-power: Low power supply low-power-recover: Low power supply Recover memusage-high: Memory usage has risen too high memusage-ok: Memory usage has fallen back to acceptable levels netusage-high: Network utilization has risen too high netusage-ok: Network utilization has fallen back to acceptable levels paging-high: Paging activity has risen too high paging-ok: Paging activity has fallen back to normal levels power-redundancy-mismatch: Power redundancy mismatch process-exit: A process in the system unexpectedly exited sm-restart: Subnet Manager restarted for parameter change sm-start: Subnet Manager started sm-stop: Subnet Manager stopped temperature-too-high: Temperature has risen too high unexpected-cluster-join: A node has unexpectedly joined the cluster unexpected-cluster-leave: A node has unexpectedly left the cluster unexpected-cluster-size: The number of nodes in the cluster is unexpected All events for which autosupport emails will be sent: liveness-failure: A process in the system was detected as hung process-crash: A process in the system has crashed</pre>
<p>Related Commands</p>	
<p>Notes</p>	

8.4 Buffer Histograms Monitoring



As it is becoming increasingly complex to manage networks, and network administrators need more tools to understand network behavior, it is necessary to provide basic information about network performance, identify network bottlenecks, and provide information for the purposes of network optimization and future planning.

Therefore, network administrators are required to constantly review network port behavior, record port buffer consumption, and identify shortage in buffer resources and record flows which lead to the excessive buffer consumption. MLNX-OS provides the following mechanisms to perform these tasks:

- Sampling (histograms)—a network administrator can enable a sampling of the port buffer occupancy, record occupancy changes over time, and provide information for different levels of buffer occupancy, and amount of time the buffer has been occupied during the observation period.

- Thresholds—thresholds may be enabled per port to record the network time when port buffer occupancy crosses the defined threshold and when buffer occupancy drops below it.
- Flow recording—a record of the most active flows which cause an excessive usage of the port buffers may be kept. Once enabled, the system may identify flow patterns and present a user with a list of flows, based on which a network administrator can rearrange distribution of the data flows in the network and minimize data loss.

8.4.1 Buffer Histograms and Thresholds Commands

8.4.1.1 protocol telemetry

	protocol telemetry no protocol telemetry Unhides telemetry config CLIs. The no form of the command hides telemetry config CLIs.
Syntax Description	N/A
Default	Hidden
Configuration Mode	config
History	3.6.3004
Example	<code>switch (config) # protocol telemetry</code>
Related Commands	
Notes	

8.4.1.2 telemetry shutdown

	telemetry shutdown no telemetry shutdown Disables the telemetry protocol, threshold detection, and histogram fetching for all sampling enabled interfaces without changing any internal configuration. The no form of the command enables telemetry protocol.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.6.3004
Example	<code>switch (config) # no telemetry shutdown</code>
Related Commands	protocol telemetry
Notes	

8.4.1.3 telemetry sampling

	interface ib <slot>/<port> telemetry sampling no interface ib <slot>/<port> telemetry sampling Enables sampling (histogram fetching) for a specific InfiniBand interface. The no form of the command disables sampling (histogram fetching).
--	--

Syntax Description	N/A
Default	N/A
Configuration Mode	config interface ib
History	3.6.3004
Example	<code>switch (config interface ib 1/1) # telemetry sampling</code>
Related Commands	protocol telemetry
Notes	

8.4.1.4 telemetry sampling log

	telemetry sampling log <time> no telemetry sampling log <time> Enables the log interval value (histogram fetching) from device. The no form of the command disables the log interval value.	
Syntax Description	time	Input range: 100-60000 (in msec)
Default	1000 millisecond	
Configuration Mode	config	
History	3.6.3004	
Example	<code>switch (config) # telemetry sampling log 1000</code>	
Related Commands	protocol telemetry	
Notes		

8.4.1.5 telemetry threshold

	telemetry threshold no telemetry threshold Enables telemetry threshold on hardware. The no form of the command disables threshold.	
Syntax Description	N/A	
Default	false	
Configuration Mode	config interface ib	
History	3.6.4006	
Example	<code>switch (config interface ib 1/1) # telemetry threshold</code>	
Related Commands		
Notes		

8.4.1.6 telemetry threshold level

	telemetry threshold level <level> no telemetry threshold level Configures threshold level in hardware per port. The no form of the command resets threshold to default value.	
Syntax Description	level	Input range: 96-1000000 (in bytes and in increments of 96)
Default	69984	
Configuration Mode	config interface ib	
History	3.6.4006	
Example	switch (config interface ib 1/1) # telemetry threshold level 288	
Related Commands		
Notes		

8.4.1.7 telemetry threshold log

	telemetry threshold log no telemetry threshold log Enables logging of threshold events in syslog. The no form of the command disables logging.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config	
History	3.6.4006	
Example	switch (config) # telemetry threshold log	
Related Commands		
Notes		

8.4.1.8 telemetry threshold record

	telemetry threshold record no telemetry threshold record Enables top talker configuration. The no form of the command disables top talker configuration.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config interface ib	
History	3.6.6105	
Example	switch (config interfaces ib 1/2) # telemetry threshold record	
Related Commands	clear telemetry threshold record show telemetry threshold record	

Notes	<ul style="list-style-type: none"> • When top talker is enabled, the minimal threshold window supported is 20 msec • Due to event timing issues, very short threshold events may not gather sufficient traffic samples to allow top-talker analysis. As a result, top-talkers may not be fully displayed in the relevant show command.
-------	--

8.4.1.9 telemetry threshold syslog

	<pre>telemetry threshold syslog <time> no telemetry threshold syslog <time></pre> <p>The command sets threshold events logging rate on per hour basis. The no form of the command sets the logging rate back to default.</p>	
Syntax Description	time	Max rate per hour Range: 1-3600
Default	100	
Configuration Mode	config	
History	3.6.4006	
Example	switch (config) # telemetry threshold syslog 400	
Related Commands		
Notes		

8.4.1.10 clear telemetry

	<pre>clear telemetry {threshold sampling} [interface ib <port-id>]</pre> <p>Clears telemetry data.</p>	
Syntax Description	port-id	InfiniBand interface ID
Default	N/A	
Configuration Mode	config interface ib	
History	3.6.5000	
Example	switch (config interface ib 1/12) # clear telemetry threshold level 288	
Related Commands		
Notes		

8.4.1.11 clear telemetry threshold

	<pre>clear telemetry threshold [interface <type> <if>]</pre> <p>Clears threshold and top talker data.</p>	
Syntax Description	type	Available values:ib
Default	N/A	
Configuration Mode	config	
History	3.6.6105	
Example	switch (config) # clear telemetry threshold interface ib 1/34-1/36	

Related Commands	
Notes	

8.4.1.12 clear telemetry threshold record

	clear telemetry threshold record [interface ib <if>] Clears top talker data.	
Syntax Description	if	InfiniBand interface ID
Default	N/A	
Configuration Mode	config	
History	3.6.6105	
Example	switch (config) # clear telemetry threshold record interface ib 1/34-1/36	
Related Commands	telemetry threshold record show telemetry threshold record	
Notes		

8.4.1.13 stats export csv telemetry

	stats export csv telemetry <slot>/<port>[/<subport>][filename <name>] [after * *] [before * *] Exports histograms collected by stats to a csv file.	
Syntax Description	slot/port	Port number
	subport	Subport number to be used if a port is split
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
	3.9.0500	Updated example
Example	switch (config) # stats export csv telemetry 1/1/4-ucast after 2020/03/16 10:54:58 before 2020/03/16 11:16:24 Generated report file: telemetry-20200316-111704.csv	
Related Commands		
Notes		

8.4.1.14 file stats telemetry delete

	file stats telemetry delete <filename> Deletes the given .csv file created by “stats export” command to user directory.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config	
History	3.6.3004	

Example	switch (config) # file stats telemetry delete telemetry-20171006-102158.csv
Related Commands	
Notes	

8.4.1.15 file stats telemetry delete latest

	file stats telemetry delete latest Delete the latest stats telemetry file.
Syntax Description	N/A
Default	N/A
Configuration Mode	Configure terminal
History	3.8.1000
Example	(config) # file stats telemetry delete latest
Related Commands	file stats telemetry delete <file_name> file stats telemetry delete all
Notes	

8.4.1.16 file stats telemetry delete all

	file stats telemetry delete all Deletes all stats telemetry files from machine.
Syntax Description	N/A
Default	N/A
Configuration Mode	Configure terminal
History	3.8.1000
Example	(config) # file stats telemetry delete all
Related Commands	file stats telemetry delete <file_name> file stats telemetry delete latest
Notes	

8.4.1.17 file stats telemetry upload

	file stats telemetry upload <filename> <upload-url> Uploads .csv file created by “stats export” command to user directory.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.3004

Example	switch (config) # file stats telemetry upload telemetry-20170119-102715.csv scp:// username:password@server//directory Password (if required): *****
Related Commands	
Notes	

8.4.1.18 file stats telemetry upload latest

	file stats telemetry upload latest <upload-url> Upload the latest stats telemetry file to a remote host.
Syntax Description	N/A
Default	N/A
Configuration Mode	Configure terminal
History	3.8.1000
Example	(config) # file stats telemetry upload latest scp://user:pass@10.135.155.8/tmp
Related Commands	file stats telemetry upload <file_name> file stats telemetry upload all
Notes	

8.4.1.19 file stats telemetry upload all

	file stats telemetry upload all <upload_url> Upload all stats telemetry files to a remote host.
Syntax Description	
Default	N/A
Configuration Mode	Configure terminal
History	3.8.1000
Example	(config) # file stats telemetry upload all scp://user:pass@10.135.155.8/tmp
Related Commands	file stats telemetry upload <file_name> file stats telemetry upload latest
Notes	

8.4.1.20 show telemetry

	show telemetry Displays the global configuration of telemetry properties.
Syntax Description	N/A
Default	N/A

Configuration Mode	config																																																		
History	3.6.4000																																																		
Example																																																			
<pre>switch (config) # show telemetry Telemetry Status : Enabled H/W Sampling Interval(nsec) : 512 S/W Sampling Interval(ms) : 1000 Threshold Logging : Disabled Threshold Logging(rate per hour) : 100</pre> <table border="1"> <thead> <tr> <th>Interface</th> <th>Sampling</th> <th>Threshold</th> <th>Record</th> <th>Level (bytes)</th> </tr> </thead> <tbody> <tr> <td>IB1/1</td> <td>Disabled</td> <td>Enabled</td> <td>Enabled</td> <td>100 (96)</td> </tr> <tr> <td>IB1/2</td> <td>Disabled</td> <td>Enabled</td> <td>Enabled</td> <td>100 (96)</td> </tr> <tr> <td>IB1/3</td> <td>Disabled</td> <td>Disabled</td> <td>Disabled</td> <td>N/A</td> </tr> <tr> <td>IB1/4</td> <td>Disabled</td> <td>Disabled</td> <td>Disabled</td> <td>N/A</td> </tr> <tr> <td>IB1/5</td> <td>Disabled</td> <td>Disabled</td> <td>Disabled</td> <td>N/A</td> </tr> <tr> <td>IB1/6</td> <td>Disabled</td> <td>Disabled</td> <td>Disabled</td> <td>N/A</td> </tr> <tr> <td>IB1/7</td> <td>Disabled</td> <td>Disabled</td> <td>Disabled</td> <td>N/A</td> </tr> <tr> <td>...</td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td>IB1/36</td> <td>Disabled</td> <td>Disabled</td> <td>Disabled</td> <td>N/A</td> </tr> </tbody> </table>		Interface	Sampling	Threshold	Record	Level (bytes)	IB1/1	Disabled	Enabled	Enabled	100 (96)	IB1/2	Disabled	Enabled	Enabled	100 (96)	IB1/3	Disabled	Disabled	Disabled	N/A	IB1/4	Disabled	Disabled	Disabled	N/A	IB1/5	Disabled	Disabled	Disabled	N/A	IB1/6	Disabled	Disabled	Disabled	N/A	IB1/7	Disabled	Disabled	Disabled	N/A	...					IB1/36	Disabled	Disabled	Disabled	N/A
Interface	Sampling	Threshold	Record	Level (bytes)																																															
IB1/1	Disabled	Enabled	Enabled	100 (96)																																															
IB1/2	Disabled	Enabled	Enabled	100 (96)																																															
IB1/3	Disabled	Disabled	Disabled	N/A																																															
IB1/4	Disabled	Disabled	Disabled	N/A																																															
IB1/5	Disabled	Disabled	Disabled	N/A																																															
IB1/6	Disabled	Disabled	Disabled	N/A																																															
IB1/7	Disabled	Disabled	Disabled	N/A																																															
...																																																			
IB1/36	Disabled	Disabled	Disabled	N/A																																															
Related Commands																																																			
Notes																																																			

8.4.1.21 show telemetry threshold record

	show telemetry threshold record [interface ib <interface-id> <interface-id-range>] Displays top talker events for all configured ports.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.4006	
	3.6.6105	Updated example
	3.6.8100	Updated example
Example		
<pre>switch (config) # show telemetry threshold record interface ib 1/11-1/12 ----- Event-id Date Time Port Level Duration(100 usec) Repeated DestQP DLID SLID Percent (%) ----- 1 07/10/18 14:00:31 IB 1/11 69984 48749.77 1 2741 29 32 62.30 2 07/10/18 14:01:47 IB 1/11 69984 63936.16 1 2745 29 32 54.55</pre>		
Related Commands	clear telemetry threshold	

Notes	<ul style="list-style-type: none"> The values displayed of the SLID, DLID, and QP fields are in decimal The command supports displaying up to 1000 threshold events. As a result, if more than 1000 thresholds configured in total, some interfaces may not be displayed. Therefore, to query thresholds for a specific interface, please use the command “show telemetry threshold interface ib <interface>”.
-------	--

8.4.1.22 show telemetry sampling interface ib

	show telemetry sampling interface ib <slot>/<port> Displays telemetry histogram samples for a specific InfiniBand interface.	
Syntax Description	slot/port	InfiniBand interface number
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
Example	<pre>switch (config) # show telemetry sampling interface ib 1/32 ----- Telemetry histogram: IB1/32 System-time Bin sizes (128 nsec tx buffer was occupied in bytes range) ----- 02/09/17 <2976 35744 68512 101280 134048 166816 199584 232352 265120 265120< 12:19:03.41948 1883 8538 7802080 0 0 0 0 0 0 12:19:04.42107 830 9001 7802670 0 0 0 0 0 0 12:19:05.42249 96 9705 7802700 0 0 0 0 0 0 12:19:06.42388 32 9035 7803434 0 0 0 0 0 0 12:19:07.42573 80 9461 7802960 0 0 0 0 0 0 12:19:08.42761 160 9302 7803040 0 0 0 0 0 0 12:19:09.42915 304 9369 7802829 0 0 0 0 0 0 12:19:10.43071 96 8906 7803500 0 0 0 0 0 0 12:19:11.43215 463 8907 7803132 0 0 0 0 0 0 12:19:12.43369 256 8571 7803675 0 0 0 0 0 0</pre>	
Related Commands		
Notes	If the requested entries are more than what the DB contains, it prints the amount in the table.	

8.4.1.23 show telemetry sampling interface ib last

	show telemetry sampling interface ib <slot>/<port> last <num_of_entries> Displays fetched unicast histogram details for an InfiniBand interface.	
Syntax Description	slot/port	InfiniBand interface number
	num_of_entries	Range: 0-1000
Default	N/A	

Configuration Mode	Any command mode
History	3.6.3004
Example	
<pre>switch (config) # show telemetry sampling interface ib 1/36 last 20 Legend: 2976 bytes - between 0 - 2976 of tx bytes buffer consumed 35744 bytes - between 2977 - 35744 of tx bytes buffer consumed ----- Telemetry histogram: IB1/36 System-time bytes range) Bin sizes (128 nsec tx buffer was occupied in ----- 02/09/17 <2976 35744 68512 101280 134048 166816 199584 232352 265120 265120< 12:19:03.41948 1883 8538 7802080 0 0 0 0 0 0 12:19:04.42107 830 9001 7802670 0 0 0 0 0 0 12:19:05.42249 96 9705 7802700 0 0 0 0 0 0 12:19:06.42388 32 9035 7803434 0 0 0 0 0 0 12:19:07.42573 80 9461 7802960 0 0 0 0 0 0</pre>	
Related Commands	
Notes	If requested entries are more than what the DB contains, it prints the amount in the table.

8.4.1.24 show files stats telemetry

	show files stats telemetry [filename] Displays all files created by the command “stats export csv telemetry”.	
Syntax Description	filename	Displays stats for the specified file
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.3004	
	3.6.8008	Updated example
Example	<pre>switch (config) # show files stats telemetry telemetry-20180527-102715.csv Hostname :test-switch Report :telemetry histogram Time lower bound(UTC) :2018/05/28 05:58:10 Time upper bound(UTC) :2018/05/28 05:58:25 Export time(UTC) :2018/05/28 06:00:06 Time lower bound :2018/05/28 08:58:10 +0300 Time upper bound :2018/05/28 08:58:25 +0300 Export time :2018/05/28 09:00:06 +0300 System version :X86_64 sys_test 2018-05-15 04:02:13 x86_64</pre>	
Related Commands	stats export csv telemetry	
Notes		

8.5 Statistics and Alarms



8.5.1 Commands

8.5.1.1 stats alarm clear

	<code>stats alarm <alarm ID> clear</code> Clears alarm state.	
Syntax Description	alarm ID	Alarms supported by the system, for example: <ul style="list-style-type: none"> • <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	
Example	<pre>switch (config) # stats alarm cpu_util_indiv clear</pre>	
Related Commands	show stats alarm	
Notes		

8.5.1.2 stats alarm enable

	<code>stats alarm <alarm-id> enable</code> <code>no stats alarm <alarm-id> enable</code> 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: <ul style="list-style-type: none"> • <code>cpu_util_indiv</code>—average CPU utilization too high: percent utilization • <code>disk_io</code>—operating System Disk I/O per second too high: kilobytes per second • <code>fs_mnt</code>—free filesystem space too low: percent of disk space free • <code>intf_util</code>—network utilization too high: bytes per second • <code>memory_pct_used</code>—too much memory in use: percent of physical memory used • <code>paging</code>—paging activity too high: page faults • <code>temperature</code>—temperature is too high: degrees
Default	The default is different per alarm-id	

Configuration Mode	config
History	3.1.0000
Example	<code>switch (config) # stats alarm cpu_util_indiv enable</code>
Related Commands	show stats alarm
Notes	

8.5.1.3 stats alarm event-repeat

	<pre>stats alarm <alarm ID> event-repeat {single while-not-cleared} no stats alarm <alarm ID> event-repeat</pre> <p>Configures repetition of events from this alarm. The no form of this command resets this parameter to its default.</p>	
Syntax Description	alarm ID	<p>Alarms supported by the system, for example:</p> <ul style="list-style-type: none"> • <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
	single	Does not repeat events: only sends one event whenever the alarm changes state.
	while-not-cleared	Repeats error events until the alarm clears.
Default	single	
Configuration Mode	config	
History	3.1.0000	
Example	<code>switch (config) # stats alarm cpu_util_indiv event-repeat single</code>	
Related Commands	show stats alarm	
Notes		

8.5.1.4 stats alarm {rising | falling}

	<pre>stats alarm <alarm ID> {rising falling} {clear-threshold error-threshold} <threshold-value></pre> <p>Configure alarms thresholds.</p>
--	--

Syntax Description	alarm ID	Alarms supported by the system, for example: <ul style="list-style-type: none"> • <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
	falling	Configures alarm for when the statistic falls too low
	rising	Configures alarm for when the statistic rises too high
	error-threshold	Sets threshold to trigger falling or rising alarm
	clear-threshold	Sets threshold to clear falling or rising alarm
	threshold-value	The desired threshold value, different per alarm
Default	Default is different per alarm-id	
Configuration Mode	config	
History	3.1.0000	
Example	<code>switch (config) # stats alarm cpu_util_indiv falling clear-threshold 10</code>	
Related Commands	show stats alarm	
Notes	Not all alarms support all four thresholds.	

8.5.1.5 stats alarm rate-limit

	<code>stats alarm <alarm ID> rate-limit {count <count-type> <count> reset window <window-type> <duration>}</code> Configures alarms rate limit.	
Syntax Description	alarm ID	Alarms supported by the system, for example: <ul style="list-style-type: none"> • <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
	count-type	Long medium, or short count (number of alarms)
	reset	Set the count and window durations to default values for this alarm
	window-type	Long medium, or short count, in seconds
Default	Short window: 5 alarms in 1 hour Medium window: 20 alarms in 1 day Long window: 50 alarms in 7 days	

Configuration Mode	config
History	3.1.0000
Example	switch (config) # stats alarm paging rate-limit window long 2000
Related Commands	show stats alarm
Notes	

8.5.1.6 stats chd clear

	stats chd <CHD ID> clear Clears CHD counters.	
Syntax Description	CHD ID	CHD supported by the system, for example: <ul style="list-style-type: none"> • cpu_util—CPU utilization: percentage of time spent • cpu_util_ave—CPU utilization average: percentage of time spent • cpu_util_day—CPU utilization average: percentage of time spent • disk_device_io_hour—storage device I/O read/write statistics for the last hour: bytes • disk_io—operating system aggregate disk I/O average (KB/sec) • fs_mnt_day—filesystem system usage average: bytes • fs_mnt_month—filesystem system usage average: bytes • fs_mnt_week—filesystem system usage average: bytes • intf_day—network interface statistics aggregation: bytes • intf_hour—network interface statistics (same as “interface” sample) • intf_util—aggregate network utilization across all interfaces • memory_day—average physical memory usage: bytes • memory_pct—average physical memory usage • paging—paging activity: page faults • paging_day—paging activity: page faults • ib_day • ib_hour
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # stats chd memory_day clear	
Related Commands	show stats chd	
Notes		

8.5.1.7 stats chd enable

	stats chd <chd-id> enable no stats chd <chd-id> enable Enables the CHD. The no form of the command disables the CHD.
--	---

Syntax Description	chd-id	<p>CHD supported by the system, for example:</p> <ul style="list-style-type: none"> • cpu_util—CPU utilization: percentage of time spent • cpu_util_ave—CPU utilization average: percentage of time spent • cpu_util_day—CPU utilization average: percentage of time spent • disk_device_io_hour—storage device I/O read/write statistics for the last hour: bytes • disk_io—operating system aggregate disk I/O average: KB/sec • fs_mnt_day—filesystem system usage average: bytes • fs_mnt_month—filesystem system usage average: bytes • fs_mnt_week—filesystem system usage average: bytes • intf_day—network interface statistics aggregation: bytes • intf_hour—network interface statistics (same as “interface” sample) • intf_util - aggregate network utilization across all interfaces • memory_day—average physical memory usage: bytes • memory_pct—average physical memory usage • paging—paging activity: page faults • paging_day—paging activity: page faults • ib_day • ib_hour
Default	Enabled	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # stats chd memory_day enable	
Related Commands	show stats chd	
Notes		

8.5.1.8 stats chd compute time

	<p>stats chd <CHD ID> compute time {interval range} <number of seconds> Sets parameters for when this CHD is computed.</p>	
Syntax Description	CHD ID	<p>Possible IDs:</p> <ul style="list-style-type: none"> • cpu_util—CPU utilization: percentage of time spent • cpu_util_ave—CPU utilization average: percentage of time spent • cpu_util_day—CPU utilization average: percentage of time spent • disk_device_io_hour—storage device I/O read/write statistics for the last hour: bytes • disk_io—operating system aggregate disk I/O average: KB/sec • fs_mnt_day—filesystem system usage average: bytes • fs_mnt_month—filesystem system usage average: bytes • fs_mnt_week—filesystem system usage average: bytes • intf_day—network interface statistics aggregation: bytes • intf_hour—network interface statistics (same as “interface” sample) • intf_util—aggregate network utilization across all interfaces • memory_day—average physical memory usage: bytes • memory_pct—average physical memory usage • paging—paging activity: page faults • paging_day—paging activity: page faults • ib_day • ib_hour

	interval	Specifies calculation interval (how often to do a new calculation) in number of seconds
	range	Specifies calculation range, in number of seconds
	number of seconds	Number of seconds
Default	Different per CHD	
Configuration Mode	config	
History	3.1.0000	
Example	<code>switch (config) # stats chd memory_day compute time interval 120</code>	
Related Commands	show stats chd	
Notes		

8.5.1.9 stats export

	<code>stats export <format> <sample-id></code> Exports collected information to a file.	
Syntax Description	memory	Memory utilization
	paging	Paging I/O
	telemetry	Telemetry histogram
	cpu_util	CPU utilization
	power	Power
Default	N/A	
Configuration Mode	config	
History	3.7.1102 3.10.1000: Updated syntax description options	
Example	<code>switch (config) # stats export csv memory</code>	
Related Commands	show stats sample	
Notes		

8.5.1.10 stats sample clear

	<code>stats sample <sample ID> clear</code> Clears sample history.
--	---

Syntax Description	sample ID	Possible sample IDs are: <ul style="list-style-type: none"> • congested • cpu_util—CPU utilization: milliseconds of time spent • disk_device_io—storage device I/O statistics • disk_io—operating system aggregate disk I/O: KB/sec • fan - Fan speed • fs_mnt_bytes—filesystem usage: bytes • fs_mnt_inodes—filesystem usage: inodes • interface—network interface statistics • intf_util—network interface utilization: bytes • memory—system memory utilization: bytes • paging—paging activity: page faults • power—power supply usage • power-consumption • temperature—modules temperature • ib
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # stats sample temperature clear	
Related Commands	show stats sample	
Notes		

8.5.1.11 stats sample enable

	stats sample <sample-id> enable no states sample <sample-id> enable Enables the sample. The no form of the command disables the sample.	
Syntax Description	sample-id	Possible sample IDs are: <ul style="list-style-type: none"> • congested • cpu_util—CPU utilization: milliseconds of time spent • disk_device_io—storage device I/O statistics • disk_io—operating system aggregate disk I/O: KB/sec • fan—fan speed • fs_mnt_bytes—filesystem usage: bytes • fs_mnt_inodes—filesystem usage: inodes • interface—network interface statistics • intf_util—network interface utilization: bytes • memory—system memory utilization: bytes • paging—paging activity: page faults • power—power supply usage • power-consumption • temperature—modules temperature • ib
Default	Enabled	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # stats sample temperature enable	
Related Commands	show stats sample	

Notes	
-------	--

8.5.1.12 stats sample interval

	<code>stats sample <sample-id> interval [<interval>]</code> <code>no stats sample <sample-id> interval [<interval>]</code> Sets the sampling interval between taking of sample records. The no form of the command sets interval to default value.	
Syntax Description	sample-id	Sample name for which report file should be generated. <ul style="list-style-type: none"> • congested • cpu_util—CPU utilization: milliseconds of time spent • disk_device_io—storage device I/O statistics • disk_io—operating system aggregate disk I/O: KB/sec • fan—fan speed • fs_mnt_bytes—filesystem usage: bytes • fs_mnt_inodes—filesystem usage: inodes • interface—network interface statistics • intf_util—network interface utilization: bytes • memory—system memory utilization: bytes • paging—paging activity: page faults • power—power supply usage • power-consumption • temperature—modules temperature • ib
	interval	Measured in seconds. Range: 1 - 86400 (24 hours)
Default	Default for “interface” samples is 60 seconds	
Configuration Mode	config	
History	3.7.1102	
Example	<pre>switch (config) # stats sample interface-ethernet interval 1</pre>	
Related Commands	show stats sample	
Notes		

8.5.1.13 stats sample max-entries

	<code>stats sample <sample-id> max-entries [<max-entries>]</code> <code>no stats sample <sample-id> max-entries [<max-entries>]</code> Sets number of records to be kept in memory for the counter. The no form of the command resets the value to its default.
--	--

Syntax Description	sample-id	Sample name for which report file should be generated. <ul style="list-style-type: none"> • congested • cpu_util—CPU utilization: milliseconds of time spent • disk_device_io—storage device I/O statistics • disk_io—operating system aggregate disk I/O: KB/sec • fan—fan speed • fs_mnt_bytes—filesystem usage: bytes • fs_mnt_inodes—filesystem usage: inodes • interface—network interface statistics • intf_util—network interface utilization: bytes • memory—system memory utilization: bytes • paging—paging activity: page faults • power—power supply usage • power-consumption • temperature—modules temperature • ib
	max-entries	Number of records Range: 1-1000
Default	Default “interface” samples is 100 records	
Configuration Mode	config	
History	3.7.1102	
Example	<code>switch (config) # stats sample interface-ethernet max-entries 1000</code>	
Related Commands	show stats sample	
Notes	<ul style="list-style-type: none"> • Setting a new value will delete all sample history. • History does not persist after reboot. 	

8.5.1.14 stats clear-all

	stats clear-all Clears data for all samples, CHDs, and status for all alarms.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Example	<code>switch (config) # stats clear-all</code>
Related Commands	show stats sample
Notes	

8.5.1.15 show stats alarm

	show stats alarm [<alarm-id> [rate-limit]] Displays status of all alarms or the specified alarm.
--	---

Syntax Description	alarm-id	Available values: <ul style="list-style-type: none"> • cpu_util_indiv—average CPU utilization too high: percent utilization • disk_io—operating System Disk I/O per second too high: kilobytes per second • fs_mnt—free filesystem space too low: percent of disk space free • intf_util—network utilization too high: bytes per second • memory_pct_used—too much memory in use: percent of physical memory used • paging—paging activity too high: page faults • temperature—temperature is too high: degrees
	rate-limit	Displays rate limit parameters.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example	<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): (disabled) Alarm fs_mnt (Free filesystem space too low): ok Alarm intf_util (Network utilization too high): (disabled) Alarm memory_pct_used (Too much memory in use): (disabled) Alarm paging (Paging activity too high): ok Alarm temperature (Temperature is too high): ok</pre>	
Related Commands	stats alarm	
Notes		

8.5.1.16 show stats chd

	show stats chd [<chd-id>] Displays configuration of all statistics CHDs.	
Syntax Description	chd-id	Available values: <ul style="list-style-type: none"> • cpu_util_indiv—average CPU utilization too high: percent utilization • disk_io—operating System Disk I/O per second too high: kilobytes per second • fs_mnt—free filesystem space too low: percent of disk space free • intf_util—network utilization too high: bytes per second • memory_pct_used—too much memory in use: percent of physical memory used • paging—paging activity too high: page faults • temperature—temperature is too high: degrees
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	

Example	<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)</pre>
Related Commands	stats chd
Notes	

8.5.1.17 show stats cpu

	<pre>show stats cpu</pre> <p>Displays some basic stats about CPU utilization:</p> <ul style="list-style-type: none"> • the current level • the peak over the past hour • the average over the past hour
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<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%</pre>
Related Commands	
Notes	

8.5.1.18 show stats sample

	<pre>show stats sample [<sample-id>]</pre> <p>Displays sampling interval for all samples, or the specified one.</p>	
Syntax Description	sample-id	<p>Sample name for which report file should be generated.</p> <ul style="list-style-type: none"> • congested • cpu_util—CPU utilization: milliseconds of time spent • disk_device_io—storage device I/O statistics • disk_io—operating system aggregate disk I/O: KB/sec • fan—fan speed • fs_mnt_bytes—filesystem usage: bytes • fs_mnt_inodes—filesystem usage: inodes • interface—network interface statistics • intf_util—network interface utilization: bytes • memory—system memory utilization: bytes • paging—paging activity: page faults • power—power supply usage • power-consumption • temperature—modules temperature • ib

Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show stats sample fan Sample "fan" (Fan speed): Enabled: yes Sampling interval: 1 minute 11 seconds</pre>
Related Commands	
Notes	

8.5.1.19 show stats sample data

	<pre>show stats sample <sample-id> data [interface {ethernet port-channel mlag-port-channel} <device/port> [counter <counter-name>]] [group name <group-name> [counter <counter-name>]] [max-samples {<max-samples> all}]</pre> <p>Displays history of counter values (i.e., collected information for a sample).</p>	
Syntax Description	sample-id	<p>Sample name for which report file should be generated.</p> <ul style="list-style-type: none"> • congested • cpu_util—CPU utilization: milliseconds of time spent • disk_device_io—storage device I/O statistics • disk_io—operating system aggregate disk I/O: KB/sec • fan—fan speed • fs_mnt_bytes—filesystem usage: bytes • fs_mnt_inodes—filesystem usage: inodes • interface—network interface statistics • intf_util—network interface utilization: bytes • memory—system memory utilization: bytes • paging—paging activity: page faults • power—power supply usage • power-consumption • temperature—modules temperature • ib
	interface	Allows limiting output to a particular interface's counters
	group	Allows limiting output to a particular group of counters
	counter	Allows limiting output to a particular counter. This option is available only if the option interface or group is chosen.
	max-samples	Allows choosing a number of counter records to display. Range: 1-1000 records. The "all" option is meant for all available records. By default, 20 counter records are displayed.
Default	N/A	
Configuration Mode	Any command mode	
History	<p>3.7.1102</p> <p>3.8.1000: Modified configuration mode & example</p> <p>3.9.2000: Modified note and example</p>	
Example		


```
switch (config) # show stats sample interface-ethernet data interface ethernet 1/1 max-samples 1
Sampling data for Interface ethernet counters:
Eth1/1:
```

Name	Timestamp	Value
Rx_packets	2000/12/25 10:27:53	0
Rx_unicast_packets	2000/12/25 10:27:53	0
Rx_multicast_packets	2000/12/25 10:27:53	0
Rx_broadcast_packets	2000/12/25 10:27:53	0
Rx_bytes	2000/12/25 10:27:53	0
Rx_discard_packets	2000/12/25 10:27:53	0
Rx_error_packets	2000/12/25 10:27:53	0
Rx_fcs_errors	2000/12/25 10:27:53	0
Rx_undersize_packets	2000/12/25 10:27:53	0
Rx_oversize_packets	2000/12/25 10:27:53	0
Rx_pause_packets	2000/12/25 10:27:53	0
Rx_unknown_control_opcode	2000/12/25 10:27:53	0
Rx_symbol_errors	2000/12/25 10:27:53	0
Rx_packets_of_64_bytes	2000/12/25 10:27:53	0
Rx_packets_of_65-127_bytes	2000/12/25 10:27:53	0
Rx_packets_of_128-255_bytes	2000/12/25 10:27:53	0
Rx_packets_of_256-511_bytes	2000/12/25 10:27:53	0
Rx_packets_of_512-1023_bytes	2000/12/25 10:27:53	0
Rx_packets_of_1024-1518_bytes	2000/12/25 10:27:53	0
Rx_packets_Jumbo	2000/12/25 10:27:53	0
Tx_packets	2000/12/25 10:27:53	0
Tx_unicast_packets	2000/12/25 10:27:53	0
Tx_multicast_packets	2000/12/25 10:27:53	0
Tx_broadcast_packets	2000/12/25 10:27:53	0
Tx_bytes	2000/12/25 10:27:53	0
Tx_discard_packets	2000/12/25 10:27:53	0
Tx_error_packets	2000/12/25 10:27:53	0
Tx_hoq_discard_packets	2000/12/25 10:27:53	0
Tx_pause_packets	2000/12/25 10:27:53	0
Tx_pause_duration	2000/12/25 10:27:53	0

Related Commands

Notes

- Filtering keyword depends on chosen <sample-id>.
- Notice that this is a history of counters. Autocompletion and output can contain information for groups (interfaces) that is not present anymore in the system, and vice versa. If counters are not sampled, they will not appear in the output.
- Output of collected information is implemented only for the following samples:
 - memory
 - paging
 - power

8.6 Management Information Bases (MIBs)



8.6.1 Calculating of entPhysicalIndex in the Entity MIB

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

1. Module layer—includes modules located on system (e.g., cables, fan, power supply, and so forth). See the [module type breakdown table](#) for more details.
2. Device layer—a number identifying the specific device that is associated with the module (e.g., ASIC on a leaf, fan on the management board, and so forth).
3. Sensor layer—a number identifying the specific sensor that is associated with the device (e.g., fan sensors, temperature sensors, power sensors, and so forth).

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

The physical entities in the system (other than port modules) use the following index schema:

Mod. Type ID	Module Index		Device Identifier				Sensor Type and Index	
1	2	3	4	5	6	7	8	9
Layer 1			Layer 2				Layer 3	

Quantum systems use the following index schema for port modules and port module sensors:

Mod. Type ID	Port Module Identifier							Port module Sensor index TX sensors in range 1..39 RX sensors in range 41..79	
1	2	3	4	5	6	7	8	9	10
Layer 1	Layer 2							Layer 3	

Switch-IB, Switch-IB 2 use the following index schema for port modules and port module sensors:

Mod. Type ID	Port Module Identifier					Port Module Sensor Type 0 for TX 1 for RX	Sensor index	
1	2	3	4	5	6	7	8	9
Layer 1	Layer 2					Layer 3		

Module type breakdown:

Number	Description
1	Chassis
2	Management
3	Spine
4	Leaf
5	Fan
6	Power supply
7	BBU
8	x86 CPU
9	Port module
Physical entities—10 digits representation	
1	Port module

Port module 9 digits representation is kept for backwards compatibility.

8.6.2 Examples

- entPhysicalIndex with value 401191311
 - 9 digits representation.
 - Layer 1 is “401”—“4” indicates a leaf (see [module type breakdown table](#)) and “01” indicates leaf at index #1 (i.e., leaf 01)
 - Layer 2 is “1913”—this is the identifier for one of the QSFP-ASIC in the system
 - Layer 3 is “11”—this is the identifier for temperature sensor #1
 - The description for this physical entity (appears in entPhysicalDescr column of the MIB) would be: L01/QSFP-ASIC-1/T1
- entPhysicalIndex with value 501020021
 - 9 digits representation.
 - Layer 1 is “501”—“5” indicates a fan (see [module type breakdown table](#)) and “01” indicates fan at index #1 (i.e., fan 01)
 - Layer 2 is “0200”—this is the identifier for general fan in the system
 - Layer 3 is “21”—this is the identifier for fan sensor #1
 - The description for this physical entity (appears in entPhysicalDescr column of the MIB) would be: FAN1/FAN/F1
- For entPhysicalIndex with value 1000012700
 - 10 digits representation.
 - Layer 1 is “1”—port module (see [module type breakdown table](#)).
 - Layer 2 is “127”—port identifier
 - Layer 3 is “00”—no sensors for this port module
- For entPhysicalIndex with value 1000012742
 - 10 digits representation.
 - Layer 1 is “1”—port module (see [module type breakdown table](#)).
 - Layer 2 is “127”—port identifier
 - Layer 3 is “42”—sensor in the range 41..79 indicates an RX sensor

9 User Management, Authentication, & Security

- [User Management & Security](#)
- [Cryptographic \(X.509, IPSec\) and Encryption](#)

9.1 User Management & Security



9.1.1 User Accounts

There are two general user account types: admin and monitor. As admin, the user is privileged to execute all the available operations. As monitor, the user can execute operations that display system configuration and status, or set terminal settings.

User Role	Default Password
admin	admin
monitor	monitor

9.1.2 Authentication, Authorization, and Accounting (AAA)

AAA is a term describing a framework for intelligently controlling access to computer resources, enforcing policies, auditing usage, and providing the information necessary to bill for services. These combined processes are considered important for effective network management and security. The AAA feature allows you to verify the identity of, grant access to, and track the actions of users managing the system. The Remote Access Dial-In User Service (RADIUS) or Terminal Access Controller Access Control device Plus (TACACS+) or Lightweight Directory Access Protocol (LDAP) protocols are supported by the MLNX-OS switch.

- **Authentication**—authentication provides the initial method of identifying each individual user, typically by entering a valid username and password before access is granted. The AAA server compares a user's authentication credentials with the user credentials stored in a database. If the credentials match, the user is granted access to the network or devices. If the credentials do not match, authentication fails and network access is denied.
- **Authorization**—following the authentication, a user must gain authorization for performing certain tasks. After logging into a system, for instance, the user may try to issue commands. The authorization process determines whether the user has the authority to issue such commands. Simply put, authorization is the process of enforcing policies: determining what types or qualities of activities, resources, or services a user is permitted. Usually, authorization occurs within the context of authentication. Once you have authenticated a user, they may be authorized for different types of access or activity.
- **Accounting**—the last level is accounting, which measures the resources a user consumes during access. This includes the amount of system time or the amount of data a user has sent and/or received during a session. Accounting is carried out by logging of session statistics and

usage information, and is used for authorization control, billing, trend analysis, resource utilization, and capacity planning activities.

Authentication, authorization, and accounting services are often provided by a dedicated AAA server, a program that performs these functions. Network access servers interface with AAA servers using the Remote Authentication Dial-In User Service (RADIUS) protocol.

9.1.3 User Re-authentication

Re-authentication prevents users from accessing resources or perform tasks for which they do not have authorization. If credential information (e.g., AAA server information like IP address, key, port number, and so forth) that has been previously used to authenticate a user is modified, that user gets immediately logged out and then asked to re-authenticate.

9.1.4 RADIUS

RADIUS (Remote Authentication Dial-In User Service), widely used in network environments, is a client/server protocol and software that enables remote access servers to communicate with a central server to authenticate dial-in users and authorize their access to the requested system or service. It is commonly used for embedded network devices such as routers, modem servers, switches and so on. RADIUS is currently the de-facto standard for remote authentication. It is prevalent in both new and legacy systems.

It is used for several reasons:

- RADIUS facilitates centralized user administration
- RADIUS consistently provides some level of protection against an active attacker

9.1.5 TACACS+

TACACS (Terminal Access Controller Access Control System), widely used in network environments, is a client/server protocol that enables remote access servers to communicate with a central server to authenticate dial-in users and authorize their access to the requested system or service. It is commonly used for providing NAS (Network Access Security). NAS ensures secure access from remotely connected users. TACACS implements the TACACS Client and provides the AAA (Authentication, Authorization, and Accounting) functionalities.

TACACS is used for several reasons:

- Facilitates centralized user administration
- Uses TCP for transport to ensure reliable delivery
- Supports inbound authentication, outbound authentication and change password request for the authentication service
- Provides some level of protection against an active attacker

9.1.6 LDAP

LDAP (Lightweight Directory Access Protocol) is an authentication protocol that allows a remote access server to forward a user's log-on password to an authentication server to determine whether access can be allowed to a given system. LDAP is based on a client/server model. The switch acts as

a client to the LDAP server. A remote user (the remote administrator) interacts only with the switch, not with the back-end server and database.

LDAP authentication consists of the following components:

- A protocol with a frame format that utilizes TCP over IP
- A centralized server that stores all the user authorization information
- A client: in this case, the switch

Each entry in the LDAP server is referenced by its Distinguished Name (DN). The DN consists of the user-account name concatenated with the LDAP domain name. The following is an example DN where the the user-account name is John:

```
uid=John,ou=people,dc=domain,dc=com
```

LDAP supports user membership in groups. If remote user is a member of admin or monitor group, it will be logged with admin or monitor capabilities respectively.

Supported group names for mapping are as follows:

- admin
- monitor

Supported group types (objectClass) on LDAP server side are as follows:

- groupOfNames
- posixGroup

9.1.7 System Secure Mode

System secure mode is a state that configures the switch system to run secure algorithms in compliance with FIPS 140-2 requirements. In this mode, unsecure algorithms are disabled and unsecure feature configurations are disallowed.

In this mode the system supports Federal Information Processing Standards (FIPS) 140-2, Security Requirements for Cryptographic Modules, which is a NIST (National Institute of Standards and Technology) publication that specifies the requirement for system cypher functionality.

When this mode is activated, all the modules which are used by the system are verified to work in compliance with the secure mode.

Note that if system fails to load in secure mode it is loaded in non-secure mode.

Prerequisites:

1. Disable SNMPv1 and v2.

```
switch (config) # no snmp-server enable communities
```

2. Only allow SNMPv3 users with sha and aes-128.

```
switch (config) # snmp-server user <username> v3 auth sha <password1> priv aes-128 <password2>
```

3. Only allow SNMPv3 traps with sha and aes-128.

```
switch (config) # snmp-server host <ip-address> informs version 3 user <username> auth sha <password1> priv aes-128 <password2>
```

4. Only allow SSHv2.

```
switch (config) # ssh server min-version 2
```

5. Enable SSH server strict security mode.

```
switch (config) # ssh server security strict
```

6. Disable HTTP access.

```
switch (config) # no web http enable
```

7. Enable HTTPS strict cyphers.

```
switch (config) # web https ssl ciphers TLS1.2
```

If a necessary prerequisite is not fulfilled the system does not activate secure mode and issues an advisory message accordingly.

Secure mode is not supported on modular switch systems.

To activate secure mode, do the following:

```
switch (config) # system secure-mode enable
Warning! Configuration is about to be saved and the system will be reloaded.
Type 'YES' to confirm the change in secure mode: YES
```

To deactivate secure mode, do the following:

```
switch (config) # no system secure-mode enable
Warning! Configuration is about to be saved and the system will be reloaded.
Type 'YES' to confirm the change in secure mode: YES
```

To verify secure mode configuration and state, do the following:

```
switch (config)# show system secure-mode
Secure mode configured: yes
Secure mode enabled: yes
```

9.1.8 User Management and Security Commands



- [9.1.8.1 User Accounts](#)
 - [9.1.8.1.1 username](#)
 - [9.1.8.1.2 show usernames](#)
 - [9.1.8.1.3 show users](#)
 - [9.1.8.1.4 show whoami](#)
 - [9.1.8.1.5 password](#)
 - [9.1.8.1.6 show password hardening](#)
- [9.1.8.2 AAA Methods](#)
 - [9.1.8.2.1 aaa accounting](#)
 - [9.1.8.2.2 aaa authentication login](#)
 - [9.1.8.2.3 aaa authentication attempts fail-delay](#)
 - [9.1.8.2.4 aaa authentication attempts track](#)
 - [9.1.8.2.5 aaa authentication attempts lockout](#)
 - [9.1.8.2.6 aaa authentication attempts class-override](#)
 - [9.1.8.2.7 aaa authentication attempts reset](#)
 - [9.1.8.2.8 clear aaa authentication attempts](#)
 - [9.1.8.2.9 aaa authorization](#)
 - [9.1.8.2.10 show aaa](#)
 - [9.1.8.2.11 show aaa authentication attempts](#)
- [9.1.8.3 RADIUS](#)
 - [9.1.8.3.1 radius-server](#)
 - [9.1.8.3.2 radius-server host](#)
 - [9.1.8.3.3 show radius](#)
- [9.1.8.4 TACACS+](#)
 - [9.1.8.4.1 tacacs-server](#)
 - [9.1.8.4.2 tacacs-server host](#)
 - [9.1.8.4.3 show tacacs](#)
- [9.1.8.5 LDAP](#)
 - [9.1.8.5.1 ldap enable](#)
 - [9.1.8.5.2 ldap base-dn](#)
 - [9.1.8.5.3 ldap bind-dn/bind-password](#)
 - [9.1.8.5.4 ldap group-attribute/group-dn](#)
 - [9.1.8.5.5 ldap nested-group-search](#)
 - [9.1.8.5.6 ldap nested-group-depth](#)
 - [9.1.8.5.7 ldap nested-group-count](#)
 - [9.1.8.5.8 ldap host](#)
 - [9.1.8.5.9 ldap hostname-check enable](#)
 - [9.1.8.5.10 ldap login-attribute](#)

- [9.1.8.5.11 ldap port](#)
- [9.1.8.5.12 ldap referrals](#)
- [9.1.8.5.13 ldap scope](#)
- [9.1.8.5.14 ldap ssl](#)
- [9.1.8.5.15 ldap timeout](#)
- [9.1.8.5.16 ldap version](#)
- [9.1.8.5.17 show ldap](#)
- [9.1.8.5.18 show ldap crl](#)
- [9.1.8.6 System Secure Mode](#)
 - [9.1.8.6.1 system secure-mode enable](#)
 - [9.1.8.6.2 show system secure-mode](#)
 - [9.1.8.6.3 show secure-boot-status](#)

9.1.8.1 User Accounts

9.1.8.1.1 username

	<p>username <username> [capability <cap> disable [login password] disconnect full-name <name> nopassword password [0 7] <password>] no username <username> [capability disable [login password] full-name] Creates a user and sets its capabilities, password and name. The no form of the command deletes the user configuration.</p>	
Syntax Description	username	<p>Specifies a username and creates a user account. New users are created initially with admin privileges but is disabled. Allowed characters for the username:</p> <ul style="list-style-type: none"> • a-z • A-Z • 0-9 • period (.), underscore (_), hyphen (-) <p>Any single character or combination of characters from the above is allowed except for a period "." in a single form.</p>
	capability <cap>	<p>Defines user capabilities.</p> <ul style="list-style-type: none"> • 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]	<ul style="list-style-type: none"> • Disable—disable this account • Disable login—disable all logins to this account • Disable password—disable login to this account using a local password
	disconnect	Logs out the specified user from the system.
	name	Full name of the user.
	nopassword	The next login of the user will not require password.
	0 7	<ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • admin • monitor 														
Configuration Mode	config														
History	<table border="1"> <tr> <td>3.1.0000</td> <td></td> </tr> <tr> <td>3.4.0000</td> <td>Updated example</td> </tr> <tr> <td>3.4.1100</td> <td>Updated example</td> </tr> <tr> <td>3.6.2002</td> <td>Added “disconnect” parameter</td> </tr> <tr> <td>3.8.1000</td> <td>Added "username" syntax description (allowed characters)</td> </tr> <tr> <td>3.8.2000</td> <td>Removed xmladmin and xmluser usernames due to XML depreciation</td> </tr> <tr> <td>3.9.0900</td> <td>Added note</td> </tr> </table>	3.1.0000		3.4.0000	Updated example	3.4.1100	Updated example	3.6.2002	Added “disconnect” parameter	3.8.1000	Added "username" syntax description (allowed characters)	3.8.2000	Removed xmladmin and xmluser usernames due to XML depreciation	3.9.0900	Added note
3.1.0000															
3.4.0000	Updated example														
3.4.1100	Updated example														
3.6.2002	Added “disconnect” parameter														
3.8.1000	Added "username" syntax description (allowed characters)														
3.8.2000	Removed xmladmin and xmluser usernames due to XML depreciation														
3.9.0900	Added note														
Example	<code>switch (config) # username monitor full-name smith</code>														
Related Commands	<code>show usernames</code> <code>show users</code>														
Notes	<ul style="list-style-type: none"> • To enable a user account, just set a password on it (or use the command “username <user> nopassword” to enable it with no password required for login) • Removing a user account does not terminate any current sessions that user has open; it just prevents new sessions from being established • Encrypted password is useful for the command “show configuration”, since the cleartext password cannot be recovered after it is set • The command "username <user> password <password>" or "username <user> password 0 <password>" are not security and will leave clear text in user's terminal (log and command history will be treated as sensitive information without clear text password). They are recommended to be replaced as "username <user> password" or "username <user> password" commands. 														

9.1.8.1.2 show usernames

	<code>show usernames</code> Displays list of users and their capabilities.						
Syntax Description	N/A						
Default	N/A						
Configuration Mode	Any command mode						
History	<table border="1"> <tr> <td>3.1.0000</td> <td></td> </tr> <tr> <td>3.8.1000</td> <td>Updated example output</td> </tr> <tr> <td>3.8.2000</td> <td>Updated example output</td> </tr> </table>	3.1.0000		3.8.1000	Updated example output	3.8.2000	Updated example output
3.1.0000							
3.8.1000	Updated example output						
3.8.2000	Updated example output						
Example							

switch (config) # show usernames	
USERNAME	FULL NAME CAPABILITY ACCOUNT STATUS
USERID	System Administrator admin Local password login disabled
admin	System Administrator admin No password required for login
monitor	System Monitor monitor Password set (SHA512)
root	Root User admin No password required for login
Related Commands	username show users
Notes	

9.1.8.1.3 show users

	show users [history] Displays logged in users and related information such as idle time and what host they have connected from.	
Syntax Description	history	Displays current and historical sessions.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example	<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) #s how users history admin pts/3 172.22.237.34 Wed Feb 1 11:56 still logged in admin pts/3 172.22.237.34 Wed Feb 1 11:42 - 11:46 (00:04) wtmp begins Wed Feb 1 11:38:10 2012</pre>	
Related Commands	username show usernames	
Notes		

9.1.8.1.4 show whoami

	show whoami Displays username and capabilities of user currently logged in.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example	<pre>switch (config) # show whoami Current user: admin Capabilities: admin</pre>	
Related Commands	username show usernames show users	

Notes	
-------	--

9.1.8.1.5 password

	password [age expiration <days> age warning <days> history < length > length minimal <length> length maximal < length > username-password-match enable complexity-class <char class> hardening enable] Configures restrictions for new passwords.	
Syntax Description	age expiration <days>	Specifies validity period of any password configured. Range: 0-365 days (0=password will not expire) Default: 365 days
	age warning <days>	Specifies how many days before expiration a warning message should be printed while logging in. Range: 0-30 days (0 indicates that a warning message will not be printed) Default: 15 days
	history < length >	Specifies how many passwords are saved per user. New password will be compared to previous passwords and will not be allowed if it is the same as an old one. Range: 0-20 passwords Default: 5 passwords
	length minimal <length>	Specifies minimal length of allowed password. Range: 1-32 characters Default: 8 characters
	length maximal < length>	Specifies maximal length of allowed password. Range: 64-80 characters Default: 64 characters
	username-password-match enable	Restricts user from having password identical to its username. Default: enabled The no form of this command will allow this.
	complexity-class <char class>	Specifies what characters must be used while configuring password. 1. none—no restrictions 2. lower 3. lower-upper 4. lower-upper-digit 5. lower-upper-digit-special Special characters allowed are: `~!@#\$%^&*()-_+=[]{};':<.> Default: lower-upper-digit
	hardening enable	Enable password restrictions. If enabled, all the above will be checked upon every new password that is being configured. Password that does not meet the requirements will be rejected. The no form will disable any password restrictions and every password will be allowed.
Default	Enabled. After upgrade, the feature will be disabled by default.	
Configuration Mode	Config	
History	3.9.2000	
Example	switch (config) # password hardening enable	
Related Commands	show password hardening	
Notes		

9.1.8.1.6 show password hardening

	show password hardening Displays all the configured password restrictions settings.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.9.2000
Example	<pre>switch (config) # show password hardening Password settings: Password hardening : enabled Min password length : 8 (characters) Max password length : 64 (characters) Character class : Lowercase, uppercase and digits Password history length : 5 Different username and password: yes Password aging : enabled Expiration warning message : 15 (days) Password age : 365 (days) switch (config) # show password hardening Password settings: Password hardening : disabled</pre>
Related Commands	password
Notes	<ul style="list-style-type: none"> • Wizard will prompt for enabling/disabling password hardening • Configuring password 7 while password hardening is enabled, will disable it

9.1.8.2 AAA Methods

9.1.8.2.1 aaa accounting

	aaa accounting changes default stop-only tacacs+ no aaa accounting changes default stop-only tacacs+ Enables logging of system changes to an AAA accounting server. The no form of the command disables the accounting.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Example	switch (config) # aaa accounting changes default stop-only tacacs+
Related Commands	show aaa

Notes	<ul style="list-style-type: none"> • TACACS+ is presently the only accounting service method supported • Change accounting covers both configuration changes and system actions that are visible under audit logging, however this feature operates independently of audit logging, so it is unaffected by the commands “logging level audit mgmt” or “configuration audit” • Configured TACACS+ servers are contacted in the order in which they appear in the configuration until one accepts the accounting data, or the server list is exhausted • Despite the name of the “stop-only” keyword, which indicates that this feature logs a TACACS+ accounting “stop” message, and in contrast to configuration change accounting, which happens after configuration database changes, system actions are logged when the action is started, not when the action has completed
-------	---

9.1.8.2.2 aaa authentication login

	aaa authentication login default <auth method> [<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	<ul style="list-style-type: none"> • local • radius • tacacs+ • ldap
Default	local	
Configuration Mode	Any command mode	
History	3.1.0000 3.7.1102—Updated notes	
Example	switch (config) # aaa authentication login default radius tacacs+ ldap local	
Related Commands	show aaa	
Notes	<ul style="list-style-type: none"> • The order in which the methods are specified is the order in which the authentication is attempted. It is recommended that “local” is one of the methods selected. • When defining a remote server that to authenticate users against, once a connection is established with it, it does not go through other authentication methods. Meaning, if local is defined first, it will not go to other methods. If a remote server is defined first and then local (radius → local), then if the radius server is reachable, the response from this server will dictate whether the switch can be accessed or not (regardless of whether the user exists on any other authentication method). 	

9.1.8.2.3 aaa authentication attempts fail-delay

	aaa authentication attempts fail-delay <time> no aaa authentication attempts fail-delay Configures delay for a specific period of time after every authentication failure. The no form of the command resets the fail-delay to its default value.	
Syntax Description	time	Range: 0-60 seconds
Default	0	

Configuration Mode	config
History	3.5.0200
Example	switch (config) # aaa authentication attempts fail-delay 1
Related Commands	
Notes	

9.1.8.2.4 aaa authentication attempts track

	aaa authentication attempts track {downcase enable} no aaa authentication attempts track {downcase enable} Configure tracking for failed authentication attempts. The no form of the command clears configuration for tracking authentication failures.	
Syntax Description	downcase	Does not convert all usernames to lowercase (for authentication failure tracking purposes only).
	enable	Disables tracking of failed authentication attempts.
Default	N/A	
Configuration Mode	config	
History	3.5.0200	
Example	switch (config) # aaa authentication attempts track enable	
Related Commands		
Notes	<ul style="list-style-type: none"> • 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. 	

9.1.8.2.5 aaa authentication attempts lockout

	aaa authentication attempts lockout {enable lock-time max-fail unlock-time} no aaa authentication attempts lockout {enable lock-time max-fail unlock-time} Configures lockout of accounts based on failed authentication attempts. The no form of the command clears configuration for lockout of accounts based on failed authentication attempts.
--	--

Syntax Description	enable	<p>Enables locking out of user accounts based on authentication failures.</p> <p>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.</p> <p>Lockouts only work if tracking is enabled. Enabling lockouts automatically enables tracking. Disabling tracking automatically disables lockouts.</p>
	lock-time	<p>Sets maximum permitted consecutive authentication failures before locking out users.</p> <p>Unlike the “max-fail” setting, this does take effect immediately for all accounts.</p> <p>If both unlock-time and lock-time are set, the unlock-time must be greater than the lock-time.</p> <p>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.</p>
	max-fail	<p>Sets maximum permitted consecutive authentication failures before locking out users.</p> <p>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.</p>
	unlock-time	<p>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.</p> <p>Unlike the “max-fail” setting, this does take effect immediately for all accounts.</p> <p>If both unlock-time and lock-time are set, the unlock-time must be greater than the lock-time.</p> <p>Careful with disabling the unlock-time, particularly if you have max-fail set to something, and have not overridden the behavior for the admin (i.e. they are subject to lockouts also). If the admin account gets locked out, and there are no other administrators who can aid, the user may be forced to boot single-user and use the pam_tallybyname command-line utility to unlock your account manually. Even if one is careful not to incur this many authentication failures, it makes the system more subject to DOS attacks.</p>
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Example	switch (config) # aaa authentication attempts logout enable	
Related Commands		
Notes		

9.1.8.2.6 aaa authentication attempts class-override

	aaa authentication attempts class-override {admin [no-lockout] unknown {no-track hash-username}} no aaa authentication attempts class-override {admin unknown {no-track hash-username}} Overrides the global settings for tracking and lockouts for a type of account. The no form of the command removes this override and lets the admin be handled according to the global settings.	
Syntax Description	admin	Overrides the global settings for tracking and lockouts for the admin account. This applies only to the single account with the username “admin”. It does not apply to any other users with administrative privileges.
	no-lockout	Prevents the admin user from being locked out though authentication failure history is still tracked (if tracking is enabled overall).
	unknown	Overrides the global settings for tracking and lockouts for unknown accounts. The “unknown” class here contains the following categories: <ul style="list-style-type: none"> • Real remote usernames which simply failed authentication • Mis-typed remote usernames • Passwords accidentally entered as usernames • Bogus usernames made up as part of an attack on the system
	hash-username	Applies a hash function to the username and stores the hashed result in lieu of the original
	no-track	Does not track authentication for such users (which of course also implies no-lockout)
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Example	<pre>switch (config) # aaa authentication attempts class-override admin no-lockout</pre>	
Related Commands		
Notes		

9.1.8.2.7 aaa authentication attempts reset

	aaa authentication attempts reset {all user <username>} [{no-clear-history no-unlock}] Clears the authentication history for and/or unlocks specified users.	
Syntax Description	all	Applies function to all users
	user	Applies function to a specific user
	no-clear-history	Leaves the history of login failures but unlocks the account
	no-unlock	Leaves the account locked but clears the history of login failures
Default	N/A	
Configuration Mode	config	
History	3.2.3000	

Example	<code>switch (config) # aaa authentication attempts reset user admin all</code>
Related Commands	
Notes	

9.1.8.2.8 clear aaa authentication attempts

	<code>clear aaa authentication attempts {all user <username>} [no-clear-history no-unlock]</code> Clears the authentication history for and/or unlocks specified users.	
Syntax Description	all	Applies function to all users.
	user	Applies function to a specific user.
	no-clear-history	Clears the history of login failures.
	no-unlock	Unlocks the account.
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Example	<code>switch (config) # aaa authentication attempts reset user admin no-clear-history</code>	
Related Commands		
Notes		

9.1.8.2.9 aaa authorization

	<code>aaa authorization map [default-user <username> order <policy> fallback]</code> <code>no aaa authorization map [default-user order fallback]</code> Sets the mapping permissions of a user in case a remote authentication is done. The no form of the command resets the attributes to default.	
Syntax Description	username	Specifies what local account the authenticated user will be logged on as when a user is authenticated (via RADIUS or TACACS+ or LDAP) and does not have a local account. If the username is local, this mapping is ignored.
	order <policy>	Sets the user mapping behavior when authenticating users via RADIUS or TACACS+ or LDAP to one of three choices. The order determines how the remote user mapping behaves. If the authenticated username is valid locally, no mapping is performed. The setting has the following three possible behaviors: <ul style="list-style-type: none"> • local-only—maps all remote users to the user specified by the command “aaa authorization map default-user <user name>”. Any vendor attributes received by an authentication server are ignored. • remote-first—if a local-user mapping attribute is returned and it is a valid local username, it maps the authenticated user to the local user specified in the attribute. Otherwise, it uses the user specified by the default-user command. • remote-only—maps a remote authenticated user if the authentication server sends a local-user mapping attribute. If the attribute does not specify a valid local user, no further mapping is tried.

	fallback	Sets the authenticating fallback behavior via RADIUS or TACACS+ or LDAP. This option attempts to authenticate username through the next authentication method listed in case of an error. <ul style="list-style-type: none"> server-err—performs fallback if an error occurs while connecting to remote AAA server (e.g., server is down, not responding, and so forth)
Default	Default user—admin Map order—remote-first Order fallback—server-err	
Configuration Mode	config	
History	3.1.0000 3.7.1000—Added “fallback” parameter 3.7.1000—Updated syntax	
Example	switch (config) # aaa authorization map default-user admin	
Related Commands	show aaa username	
Notes	<ul style="list-style-type: none"> If, for example, the user is locally defined to have admin permission, but in a remote server such as RADIUS the user is authenticated as monitor and the order is remote-first, then the user is given monitor permissions. The user must be careful when disabling AAA authorization map fallback server-err, because if the remote server stops working then the user may lock themselves out. If AAA authorization order policy is configured to remote-only, then when upgrading to 3.4.3000 or later from an older version, this policy is changed to remote-first. 	

9.1.8.2.10 show aaa

	show aaa Displays the AAA configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.7.0020—Example updated
Example	<pre>switch (config) # show aaa AAA authorization: Default User: admin Map Order: remote-first Fallback on server-err: yes Authentication method(s): local Accounting method(s): tacacs+</pre>
Related Commands	aaa accounting aaa authentication aaa authorization show aaa show usernames username
Notes	

9.1.8.2.11 show aaa authentication attempts

	show aaa authentication attempts [configured status user <username>]] Displays the current authentication, authorization and accounting settings.	
Syntax Description	authentication attempts	Displays configuration and history of authentication failures.
	configured	Displays configuration of authentication failure tracking.
	status user	Displays status of authentication failure tracking and lockouts for specific user.
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.1000 3.5.0200—Updated example	
Example		
<pre>switch (config) # show aaa authentication attempts Configuration for authentication failure tracking and locking: Track authentication failures: yes Lock accounts based on authentication failures: yes Override treatment of 'admin' user: (none) Override treatment of unknown usernames: hash-usernames Convert usernames to lowercase for tracking: no Delay after each auth failure (fail delay): none Configuration for lockouts based on authentication failures: Lock account after consecutive auth failures: 5 Allow retry on locked accounts (unlock time): after 15 second(s) Temp lock after each auth failure (lock time): none Username Known Locked Failures Last fail time Last fail from ----- - 0Q72B43EHBKT8CB5AF5PGRX3U3B3TUL4CYJP93N(*) no no 1 2020/05/20 14:29:19 ttyS0 (*) Hashed for security reasons</pre>		
Related Commands		
Notes		

9.1.8.3 RADIUS

9.1.8.3.1 radius-server

	radius-server {key <secret> retransmit <retries> timeout <seconds>} no radius-server {key retransmit timeout} Sets global RADIUS server attributes. The no form of the command resets the attributes to their default values.	
Syntax Description	secret	Sets a secret key (shared hidden text string), known to the system and to the RADIUS server.
	retries	Number of retries (0-5) before exhausting from the authentication.
	seconds	Timeout in seconds between each retry (1-60).

Default	3 seconds, 1 retry
Configuration Mode	config
History	3.1.0000
Example	switch (config) # radius-server retransmit 3
Related Commands	aaa authorization radius-server host show radius
Notes	Each RADIUS server can override those global parameters using the command “radius-server host”.

9.1.8.3.2 radius-server host

	<p>radius-server host <IP address> [enable auth-port <port> key <secret> prompt-key retransmit <retries> timeout <seconds> cipher <none eap-peap>]</p> <p>no radius-server host <IP address> [auth-port enable cipher]</p> <p>Configures RADIUS server attributes. The no form of the command resets the attributes to their default values and deletes the RADIUS server.</p>	
Syntax Description	IP address	RADIUS server IP address
	enable	Administrative enable of the RADIUS server
	auth-port	Configures authentication port to use with this RADIUS server
	port	RADIUS server UDP port number
	key	Configures shared secret to use with this RADIUS server
	prompt-key	Prompt for key, rather than entering on command line
	retransmit	Configures retransmit count to use with this RADIUS server
	retries	Number of retries (0-5) before exhausting from the authentication
	timeout	Configures timeout between each try
	seconds	Timeout in seconds between each retry (1-60)
	cipher	Configures which cipher to use for communication encryption <none eap-peap>
Default	3 seconds, 1 retry Default UDP port is 1812	
Configuration Mode	config	
History	3.1.0000 3.8.1000—Updated command description, syntax description & example	
Example	switch (config) # radius-server host fe80::202:b3ff:fe1e:8329 switch (config) # radius-server host 40.40.40.40	
Related Commands	aaa authorization radius-server show radius	
Notes	<ul style="list-style-type: none"> • RADIUS servers are tried in the order they are configured • If you do not specify a parameter for this configured RADIUS server, the configuration will be taken from the global RADIUS server configuration. Refer to the command “radius-server”. 	

9.1.8.3.3 show radius

	show radius Displays RADIUS configurations.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.6.6000—Updated example 3.8.1000—Updated command description, syntax description & example
Example	<pre>switch (config) # show radius RADIUS defaults: Key : ***** Timeout : 3 Retransmit : 1 RADIUS servers: 1.1.1.1:1812 : Enabled : yes Key : ***** Timeout : 3 (default) Retransmit : 1 (default) Cipher : none 40.40.40.40:1812: Enabled : yes Key : ***** Timeout : 3 (default) Retransmit : 1 (default)</pre>
Related Commands	aaa authorization radius-server radius-server host
Notes	

9.1.8.4 TACACS+

9.1.8.4.1 tacacs-server

	tacacs-server {key <secret> retransmit <retries> timeout <seconds>} no tacacs-server {key retransmit timeout} Sets global TACACS+ server attributes. The no form of the command resets the attributes to default values.	
Syntax Description	secret	Set a secret key (shared hidden text string), known to the system and to the TACACS+ server.
	retries	Number of retries (0-5) before exhausting from the authentication.
	seconds	Timeout in seconds between each retry. Reang: 1-60
Default	3 seconds, 1 retry	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # tacacs-server retransmit 3	

Related Commands	aaa authorization show radius show tacacs tacacs-server host
Notes	Each TACACS+ server can override those global parameters using the command "tacacs-server host".

9.1.8.4.2 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: <ul style="list-style-type: none"> • ASCII • PAP (Password Authentication Protocol)
	key	Configures shared secret to use with this TACACS+ server.
	secret	Sets a secret key (shared hidden text string), known to the system and to the TACACS+ server.
	prompt-key	Prompts for key, rather than entering key on command line.
	retransmit	Configures retransmit count to use with this TACACS+ server.
	retries	Number of retries (0-5) before exhausting from the authentication.
	timeout	Configures timeout to use with this TACACS+ server.
	seconds	Timeout in seconds between each retry. Range: 1-60
Default	3 seconds, 1 retry Default TCP port is 49 Default auth-type is PAP	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # tacacs-server host 40.40.40.40 switch (config) # tacacs-server host fe80::202:b3ff:fe1e:8329	
Related Commands	aaa authorization show tacacs tacacs-server	

Notes	<ul style="list-style-type: none"> • TACACS+ servers are tried in the order they are configured • A PAP auth-type similar to an ASCII login, except that the username and password arrive at the network access server in a PAP protocol packet instead of being typed in by the user, so the user is not prompted • If the user does not specify a parameter for this configured TACACS+ server, the configuration will be taken from the global TACACS+ server configuration. Refer to the command “tacacs-server”.
-------	--

9.1.8.4.3 show tacacs

	<pre>show tacacs</pre> Displays TACACS+ configurations.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.6.6000—Updated example
Example	<pre>TACACS+ servers: 1.1.1.1:49: Enabled : yes Auth Type : pap Key : ***** Timeout : 3 (default) Retransmit: 1 (default)</pre>
Related Commands	<pre>aaa authorization tacacs-server tacacs-server host</pre>
Notes	

9.1.8.5 LDAP

9.1.8.5.1 ldap enable

	<pre>ldap [vrf <vrf-name>] enable [force] no ldap [vrf <vrf-name>] enable</pre> Enables LDAP in VRF. The no form of the command disables LDAP in a specified VRF.		
Syntax Description	<table border="1"> <tr> <td>force</td> <td>Enables LDAP in the specified VRF while setting all relevant LDAP options to default.</td> </tr> </table>	force	Enables LDAP in the specified VRF while setting all relevant LDAP options to default.
force	Enables LDAP in the specified VRF while setting all relevant LDAP options to default.		
Default	LDAP enabled		
Configuration Mode	config		
History	3.9.2000		
Example	<pre>switch (config) # ldap vrf mgmt enable</pre>		
Related Commands			
Notes	If VRF mgmt exists, LDAP will be enabled on VRF mgmt. If there is no VRF mgmt, LDAP will be enabled on the "default" VRF.		

9.1.8.5.2 ldap base-dn

	<code>ldap base-dn <string></code> <code>no ldap base-dn</code> Sets the base distinguished name (location) of the user information in the schema of the LDAP server. The no form of the command resets the attribute to its default values.	
Syntax Description	string	A case-sensitive string that specifies the location in the LDAP hierarchy where the server should begin searching when it receives an authorization request. For example: “ou=users,dc=example,dc=com”, with no spaces. Where: <ul style="list-style-type: none"> • ou—Organizational unit • dc—Domain component • cn—Common name • sn—Surname
Default	ou=users,dc=example,dc=com	
Configuration Mode	config	
History	3.1.1000 3.4.0000—Updated example	
Example	<pre>switch (config) # ldap base-dn ou=department,dc=example,dc=com</pre>	
Related Commands	show ldap	
Notes		

9.1.8.5.3 ldap bind-dn/bind-password

	<code>ldap {bind-dn bind-password} <string></code> <code>no ldap {bind-dn bind-password}</code> 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	3.1.1000 3.4.0000—Updated example	
Example	<pre>switch (config) # ldap bind-dn my-dn switch (config) # ldap bind-password my-password</pre>	
Related Commands	show ldap	
Notes	For anonymous login, bind-dn and bind-password should be empty strings “”.	

9.1.8.5.4 ldap group-attribute/group-dn

	ldap {group-attribute {<group-att> member uniqueMember} group-dn <group-dn>} no ldap {group-attribute group-dn} Sets the distinguished name or attribute name of a group on the LDAP server. The no form of the command resets the attribute to its default values.	
Syntax Description	group-att	Specifies a custom attribute name.
	member	groupOfNames or group membership attribute.
	uniqueMember	groupOfUniqueNames membership attribute.
	group-dn	DN of group required for authorization.
Default	group-att: member group-dn: ""	
Configuration Mode	config	
History	3.1.1000 3.4.0000—Updated example	
Example	<pre>switch (config) # ldap group-attribute member switch (config) # ldap group-dn my-group-dn</pre>	
Related Commands	show ldap	
Notes	<ul style="list-style-type: none"> • 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. 	

9.1.8.5.5 ldap nested-group-search

	ldap nested-group-search no ldap nested-group-search Enable LDAP nested-group search mechanism for user-authentication group matching. The no form of the command resets the attribute to its default values.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config	
History	3.10.2000	
Example	<pre>switch (config) # ldap nested-group-search switch (config) # no ldap nested-group-search</pre>	
Related Commands	ldap nested-group-depth ldap nested-group-count show ldap	
Notes		

9.1.8.5.6 ldap nested-group-depth

	ldap nested-group-depth <1-9> no ldap nested-group-depth Sets LDAP maximum depth for nested-group search. The no form of the command resets search depth to default (3).	
Syntax Description	N/A	
Default	3	
Configuration Mode	config	
History	3.10.2000	
Example	<pre>switch (config) # ldap nested-group-depth 6 switch (config) # no ldap nested-group-depth</pre>	
Related Commands	ldap nested-group-search ldap nested-group-count show ldap	
Notes		

9.1.8.5.7 ldap nested-group-count

	ldap nested-group-count <1-10000> no ldap nested-group-count Sets LDAP maximum number of queried nested-groups. The no form of the command resets search depth to default (1000).	
Syntax Description	N/A	
Default	1000	
Configuration Mode	config	
History	3.10.2000	
Example	<pre>switch (config) # ldap nested-group-count 500 switch (config) # no ldap nested-group-count</pre>	
Related Commands	ldap nested-group-depth ldap nested-group-search show ldap	
Notes		

9.1.8.5.8 ldap host

	ldap host <ip-address> [order <number> last] no ldap host <ip-address> Adds an LDAP server to the set of servers used for authentication. The no form of the command deletes the LDAP host.	
Syntax Description	ip-address	IPv4 or IPv6 address.
	number	The order of the LDAP server.
	last	The LDAP server will be added in the last location.
Default	No hosts configured	

Configuration Mode	config
History	3.1.1000 3.4.0000—Updated example
Example	switch (config) # ldap host 10.10.10.10
Related Commands	show aaa show ldap
Notes	<ul style="list-style-type: none"> 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

9.1.8.5.9 ldap hostname-check enable

	ldap hostname-check enable no ldap hostname-check enable Enables LDAP hostname check. The no form of the command disables LDAP hostname check.
Syntax Description	N/A
Default	No hosts configured
Configuration Mode	config
History	3.6.8008
Example	switch (config) # ldap hostname-check enable
Related Commands	show aaa show ldap
Notes	

9.1.8.5.10 ldap login-attribute

	ldap login-attribute {<string> uid sAMAccountName} no ldap login-attribute Sets the attribute name which contains the login name of the user. The no form of the command resets this attribute to its default.	
Syntax Description	string	Custom attribute name.
	uid	LDAP login name is taken from the user login username.
	sAMAccountName	SAM Account name, active directory login name.
Default	sAMAccountName	
Configuration Mode	config	
History	3.1.1000 3.4.0000—Updated example	
Example	switch (config) # ldap login-attribute uid	
Related Commands	show aaa show ldap	
Notes		

9.1.8.5.11 ldap port

	ldap port <port> no ldap port Sets the TCP port on the LDAP server to connect to for authentication. The no form of the command resets this attribute to its default value.	
Syntax Description	port	TCP port number
Default	389	
Configuration Mode	config	
History	3.1.1000 3.4.0000—Updated example	
Example	<pre>switch (config) # ldap port 1111</pre>	
Related Commands	show aaa show ldap	
Notes		

9.1.8.5.12 ldap referrals

	ldap referrals no ldap referrals Enables LDAP referrals. The no form of the command disables LDAP referrals.	
Syntax Description	N/A	
Default	LDAP referrals are enabled	
Configuration Mode	config	
History	3.1.1000 3.4.0000—Updated example	
Example	<pre>switch (config) # no ldap referrals</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.	

9.1.8.5.13 ldap scope

	ldap scope <scope> no ldap scope Specifies the extent of the search in the LDAP hierarchy that the server should make when it receives an authorization request. The no form of the command resets the attribute to its default value.	
Syntax Description	scope	<ul style="list-style-type: none"> • 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	3.1.1000 3.4.0000—Updated example
Example	switch (config) # ldap scope subtree
Related Commands	show aaa show ldap
Notes	

9.1.8.5.14 ldap ssl

	ldap ssl {ca-list <options> cert-verify ciphers {all TLS1.2} crl-check {enable file fetch all [vrf <vrf-name>] <path>} mode <mode> port <port-number>} no ldap ssl {cert-verify ciphers crl-check enable mode port} Sets SSL parameter for LDAP. The no form of the command resets the attribute to its default value.	
Syntax Description	options	This command specifies the list of supplemental certificates of authority (CAs) from the certificate configuration database that is to be used by LDAP for authentication of servers when in TLS or SSL mode. The options are: <ul style="list-style-type: none"> • default-ca-list—uses default supplemental CA certificate list • none—no supplemental list, uses the built-in one only CA certificates are ignored if “ldap ssl mode” is not configured as either “tls” or “ssl”, or if “no ldap ssl cert-verify” is configured. The default-ca-list is empty in the factory default configuration. Use the command: “crypto certificate ca-list default-ca-list name” to add trusted certificates to that list. The “default-ca-list” option requires LDAP to consult the system’s configured global default CA-list for supplemental certificates.
	cert-verify	Enables verification of SSL/TLS server certificates. This may be required if the server’s certificate is self-signed, or does not match the name of the server.
	ciphers {all TLS1.2}	Sets SSL mode to be used
	crl-check enable	Enables LDAP CRL check
	crl-check file fetch	Fetches CRL from remote server. CRL must be a valid PEM file unless a proper message shown. Supported formats: SCP, HTTP, HTTPS, FTP, and FTPS.
	mode	Sets the security mode for connections to the LDAP server. <ul style="list-style-type: none"> • none—requests no encryption for the LDAP connection • ssl—the SSL-port configuration is used, an SSL connection is made before LDAP requests are sent (LDAP over SSL) • start-tls—the normal LDAP port is used, an LDAP connection is initiated, and then TLS is started on this existing connection
	vrf-name	VRF to be affected. If “vrf-name” parameter is not specified, “default” VRF will be used.
	port-number	Sets the port on the LDAP server to connect to for authentication when the SSL security mode is enabled (LDAP over SSL)

Default	cert-verify—enabled mode—none (LDAP SSL is not activated) port-number—636 ciphers—all
Configuration Mode	config
History	3.1.1000 3.2.3000—Added ca-list argument 3.4.0000—Added “ssl ciphers” parameter and Updated example 3.6.8008—Added the parameter “crl-check” 3.9.2000—Added VRF option 3.10.6000—Added note
Example	switch (config) # ldap ssl crl-check file fetch scp://root:pass@1.1.1.1/etc/pki/ crl.pem 100.0% [#####]
Related Commands	show aaa show ldap
Notes	<ul style="list-style-type: none"> • If available, the TLS mode is recommended, as it is standardized, and may also be of higher security • The port number is used only for SSL mode. If the security mode selected is TLS, the LDAP port number is used. • To use SSL option in LDAP, the server must support this accordingly. Enable SSL option on the LDAP server by editing the following: Add the following string on customer LDAP server "ldaps:///" to the SLAPD_SERVICES in the ldap config file.

9.1.8.5.15 ldap timeout

	ldap {timeout-bind timeout-search} <seconds> no ldap {timeout-bind timeout-search} Sets a global communication timeout in seconds for all LDAP servers to specify the extent of the search in the LDAP hierarchy that the server should make when it receives an authorization request. The no form of the command resets the attribute to its default value.	
Syntax Description	timeout-bind	Sets the global LDAP bind timeout for all LDAP servers.
	timeout-search	Sets the global LDAP search timeout for all LDAP servers.
	seconds	Number of seconds. Range: 1-60
Default	5 seconds	
Configuration Mode	config	
History	3.1.1000 3.4.0000—Updated example	
Example	switch (config) # ldap timeout-bind 10	
Related Commands	show aaa show ldap	
Notes		

9.1.8.5.16 ldap version

	ldap version <version> no ldap version Sets the LDAP version. The no form of the command resets the attribute to its default value.	
Syntax Description	version	Sets the LDAP version Available values: 2, 3
Default	3	
Configuration Mode	config	
History	3.1.1000 3.4.0000—Updated example	
Example	switch (config) # ldap version 3	
Related Commands	show aaa show ldap	
Notes		

9.1.8.5.17 show ldap

	show ldap Displays LDAP configurations.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.1000 3.4.0000—Updated example 3.6.8008—Updated example 3.10.2000—Updated example to reflect the following added fields: "Nested-group search," "nested-group search depth," and "nested-search maximum group count"	
Example	<pre>switch (config) # show ldap administratively : enabled VRF name : mgmt User base DN : ou=users,dc=example,dc=com User search scope : subtree Login attribute : sAMAccountName Bind DN : Bind password : ***** Group base DN : Group attribute : member Nested-group search : disabled Nested-group search depth : 3 Nested-search maximum group count : 1000 LDAP version : 3 Referrals : yes Server port : 389 Search Timeout : 5 Bind Timeout : 5 Server Hostname check : no SSL mode : none Server SSL port : 636 (not active) SSL ciphers : all (not active) SSL cert verify : yes SSL ca-list : default-ca-list SSL CRL check : no LDAP servers: No LDAP servers configured.</pre>	

Related Commands	show aaa show ldap
Notes	

9.1.8.5.18 show ldap crt

	show ldap crt Displays current CRL configured by the user.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.8008
Example	switch (config) # show ldap crt -----BEGIN CERTIFICATE----- MIIDVzCSd..... -----END CERTIFICATE-----
Related Commands	show aaa show ldap
Notes	

9.1.8.6 System Secure Mode

9.1.8.6.1 system secure-mode enable

	system secure-mode enable no system secure-mode enable Enables secure mode on the switch. The no form of the command disables secure mode.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.5.0200 3.10.2000: Added note
Example	switch (config) # system secure-mode enable Warning! Configuration is about to be saved and the system will be reloaded. Type 'YES' to confirm the change in secure mode: YES
Related Commands	user <username> password <password> ssh server min-version ssh server security strict snmp-server user no neighbor <ip-address> password ntp server disable ntp server keyID router bgp neighbor password router bgp peer-group password

Notes	<ul style="list-style-type: none"> • Before enabling secure mode, the command performs the following configuration checks: <ul style="list-style-type: none"> • NTP Key ID cannot be MD5 when secure mode is enabled • SSH min-version cannot be 1 when enabling secure mode • SSH security must be set to strict security • SNMPv3 user auth cannot be md5 when enabling secure mode • SNMPv3 user priv cannot be des when enabling secure mode • SNMPv3 trap auth cannot be md5 when enabling secure mode • SNMPv3 trap priv cannot be des when enabling secure mode • Router BGP neighbor password cannot be set when enabling secure mode • Router BGP peer-group password cannot be set when enabling with secure mode • User password hash cannot be MD5 when secure mode is enabled • Only if the check passes, secure mode is enabled on the switch system. • When secure mode is enabled extra reboot may happen after next steps: install new image and boot to newly installed image.
-------	--

9.1.8.6.2 show system secure-mode

	<pre>show system secure-mode</pre> Displays the security mode of the switch system.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.4.2300
Example	<pre>switch (config) # show system secure-mode</pre> <p>Secure mode configured: yes Secure mode enabled : yes</p>
Related Commands	system secure-mode enable
Notes	<ul style="list-style-type: none"> • “Secure mode configuration” describes the user configuration • “Secure mode enabled” describes the system state

9.1.8.6.3 show secure-boot-status

	<pre>show secure-boot-status</pre> Displays the state of the secure boot: enable or disable.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.10.1000
Example	<pre>Switch # show secure-boot-status</pre> <p>SecureBoot disabled</p>
Related Commands	
Notes	This command is only available for NDR platforms and above

9.2 Cryptographic (X.509, IPsec) and Encryption

This page contains commands for configuring, generating and modifying x.509 certificates used in the system. Certificates are used for creating a trusted SSL connection to the system.

Crypto commands also cover IPsec configuration commands used for establishing a secure connection between hosts over IP layer which is useful for transferring sensitive information.

9.2.1 System File Encryption

This feature encrypts all sensitive data on NVIDIA systems including logs certificates, keys, etc.

To activate encryption on the switch:

1. Enable encryption and configure key location as USB (if you are using a USB device).

```
switch (config)# crypto encrypt-data key-location usb key mypassword
Warning! All sensitive files are about to be encrypted
- System will perform reset factory, configuration files will be preserved
- System will be rebooted
- Active configuration will be preserved
- Do not power-off, wait for the system to boot
Type 'YES' to confirm this action: YES
```

IMPORTANT

Encryption and decryption perform “reset factory keep-config” on the switch system once configured. This means that sysdumps, logs, and images are deleted.

The key may be saved locally as well by using the parameter “local” instead of “usb” but that configuration is less secure.

2. After the system reboots, verify configuration.

```
switch (config)# show crypto encrypt-data
Sensitive files encryption:
Status:      enabled
Key location:  usb
Cipher:      aes256
```

Once encryption is enabled, reverting back to an older version while encrypted is not possible. The command “no crypto encrypt-data” must be run before attempting to downgrade to an older OS version.

If encryption is enabled, upgrading to a new OS version maintains the encryption configuration.

9.2.2 Changing a Certificate in the System

To change the default certificate for the system, to the following:

1. Import the certificate to be used (e.g., a certificate created by openssl outside the switch).

```
switch (config) # crypto certificate name <cert_name> public-cert pem "-----BEGIN CERTIFICATE-----
> MIIDYzCCAKsCCQC9EPbMuxjNBzANBgkqhkiG9w0BAQsFADBeMQswCQYDVQQGEwJJ
...
> fEt2ui9taB1d19480xDeGUxwUDX4Y0s/bQDjpp99z+cKXUe2eYzeEwnTdrCzPZuQo
> -----END CERTIFICATE-----"
Successfully installed certificate with name '<cert_name>'
```

Or use a new self-signed certificate via switch CLI and export it as a CSR (certificate signing request) and send said CSR to the root CA for signing:

```
switch (config) # crypto certificate name <cert_name> generate self-signed
Successfully generated certificate with name '<cert_name>'

switch (config) # show crypto certificate name <cert_name> csr-pem

-----BEGIN CERTIFICATE REQUEST-----
MIICuDCCAAaCAQAwczELMAkGA1UEBhMCSVMxDDAKBgNVBAGMA1RCRDEMMAoGA1UE
BwwDVEMQwwCgYDVQQKDANUQkQxDDAKBgNVBAsMA1RCRDEYMBYGA1UEAwwPYnVs
bGRvZy1xcDEtMTMzMRIwEAYJKoZIhvcNAQkBFgNUQkQwggEiMA0GCSqGSIb3DQEB
AQUAA4IBDWAwwgEKAoIABAQC34xRvH9BaBUPi1lV6kiSOAVAnOfgreWEYoWeGpWJ
XGZQBwewF4TgptYo5fZ4KcnYcQxrcW7gYycQB9Y+9vUVvvPi3b4aYc2FkoNtnC3
0BRtxEcIiwXY7LQxIA23Zuv/OlhjTkpe0+OYtpJSFeIDKMIX4Uy2BfevG06YLCaw
tuju2FLQVkeXayNK/HFLa5POpVt+16JLB1eV0bcC38Mq9JNIGPspJ7JIjo+BjzgD
43iEY41hlRzoalu78nBBd0HbAddxCP1Uc+8PLuPLCIjGbV9ehPJNWSsA/T9jUEFU
90KaI0/k05JqCXWnpvKz3opQraHsVAbsxG312pnmBTFNagMBAAGgADANBgkqhkiG
9w0BAQsFAAOCAQEAAhpgZRNW/jleyUbtGER0CzdNbJ70V8w21Gr6bDhZgrQ/I4eO
1K1D1hvrVWYRB0SSPFmCmVmFmC7BQne8xrbL2It3ZdSKd82Ts36/Uxjtb63hyt3
GBzCas7qypsbCVW42UHuD+259Yu5xpi9haspzD8Wg2ZKU5e6Sjch+JlchkM9mh/g
BQJo4shybTgPft+mFUCCyWmf5aLyQ9TrZpaUQ7cOk6BZB1RRkOVvA6uCFrwlBks
X72LleceL4fP9dtML4VMzMMaf+wOUNxWP9+1qkKMaDhroDP5q1o1r5BLs1RVet4
z7zb3xSaPrhnefoGr88WFO74d9RxLPPdHcMFw==
-----END CERTIFICATE REQUEST-----
```

2. Import key of certificate.

```
switch (config) # crypto certificate name <cert_name> private-key pem "-----BEGIN RSA PRIVATE KEY-----
> MIIIEpAIBAAKCAQEAppJnZkwbhmt71Kf/MO6cy7QmWWHhCozzWRWuWGKse+MxSmfC
...
> QAuPOVR11SyIEnYU+X0rMhc/9tgUh/8C7mBKwj7dccMmnRwz2djsjg==
> -----END RSA PRIVATE KEY-----"
```

3. Designate <cert_name> as the global default certificate for authentication of this system to clients.

```
switch (config) # crypto certificate default-cert name <cert_name>
```

4. (Optional) Import the Certificate Authority (CA) certificate which signed for the controller.

```
switch (config) # # crypto certificate name rootCA public-cert pem "-----BEGIN CERTIFICATE-----
> MIIDjzCCAnegAwIBAgIJALVou4mcQtxlMA0GCSqGSIb3DQEBwUAMF4xCzAJBgNV
...
> +ZfQIOCFs8gY4BDq73W4ugr38mqIA8UXAMPwgjCbK4NyOh0rJ1P6WT8fYzvunct
> -----END CERTIFICATE-----"
Successfully installed certificate with name 'rootCA'
```

5. Adds the “rootCA” to the default CA certificate list.

```
switch (config) # crypto certificate ca-list default-ca-list name rootCA
```

6. Save configuration.

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

7. Verify configuration.

```
switch (config) # show crypto certificate
Certificate with name 'system-self-signed'
Comment: system-generated self-signed certificate
Private Key: present
Serial Number: 0x543e2efc3a5ecdbe18b5b5e744598424
SHA-1 Fingerprint: 14e1d36035c7a5fea9f7f0f423572c9954cb9fac

Validity:
Starts: 2022/09/12 12:44:10
Expires: 2023/09/12 12:44:10
Subject:
Common Name: switch
Country: IS
State or Province: TBD
Locality: TBD
Organization: TBD
Organizational Unit: TBD
E-mail Address: TBD

Issuer:
Common Name: switch
Country: IS
State or Province: TBD
Locality: TBD
Organization: TBD
Organizational Unit: TBD
E-mail Address: TBD

Certificate with name '<cert_name>' (default-cert)
Private Key: present
Serial Number: 0xbd10f6ccbb18cd07
SHA-1 Fingerprint: 1e0e3302182ab56f2cbd3ca21722dec55299d670

Validity:
Starts: 2021/09/12 15:16:48
Expires: 2023/01/25 14:16:48
Subject:
Common Name: switch
Country: *
State or Province: Some-State
Locality: *
Organization: NVIDIA
Organizational Unit: e2e
E-mail Address: none@nowhere.com

Issuer:
Common Name: ca
Country: *
State or Province: Some-State
Locality: *
Organization: NVIDIA
Organizational Unit: e2e

Certificate with name 'rootCA'
Private Key: not present
Serial Number: 0xb568bb899c42dc65
SHA-1 Fingerprint: 9855536f6ee0177356ffbd54ffe803bc83fb4c6
Validity:
Starts: 2020/09/08 10:34:23
Expires: 2023/06/29 10:34:23

Subject:
Common Name: ca
Country: *
State or Province: Some-State
Locality: *
Organization: NVIDIA
Organizational Unit: e2e

Issuer:
Common Name: ca
Country: *
State or Province: Some-State
Locality: *
Organization: NVIDIA
Organizational Unit: e2e
```

9.2.3 Cryptographic and Encryption Commands

- [9.2.1 System File Encryption](#)
- [9.2.2 Changing a Certificate in the System](#)
- [9.2.3 Cryptographic and Encryption Commands](#)
 - [9.2.3.1 crypto encrypt-data](#)
 - [9.2.3.2 crypto ipsec ike](#)
 - [9.2.3.3 crypto ipsec peer local](#)
 - [9.2.3.4 crypto certificate ca-list](#)

- [9.2.3.5 crypto certificate default-cert](#)
- [9.2.3.6 crypto certificate generation](#)
- [9.2.3.7 crypto certificate name](#)
- [9.2.3.8 crypto certificate system-self-signed](#)
- [9.2.3.9 show crypto certificate](#)
- [9.2.3.10 show crypto encrypt-data](#)
- [9.2.3.11 show crypto ipsec](#)

9.2.3.1 crypto encrypt-data

	crypto encrypt-data key-location <local usb> key <password> no crypto encrypt-data Enables and configures system file encryption. The no form of the command decrypts sensitive information on the system.	
Syntax Description	key-location	Configures where to store the encryption key: <ul style="list-style-type: none"> • local—stores the key locally • usb—stores the key on a USB device
	key	Configures a key
Default	N/A	
Configuration Mode	config	
History	3.6.1002	
Example		
Related Commands	show crypto certificate	
Notes	<ul style="list-style-type: none"> • It is recommended to store the encryption password on a USB device rather than locally • Enabling encryption may slightly slow system performance • If the key is stored on the USB, it must be plugged into the switch in order for the switch to boot. After the switch has booted, the USB key is no longer required and, for security purposes, it is recommended to remove it after running “usb eject”. The USB key may be needed again if the switch is rebooted or if the switch needs to be decrypted. 	

9.2.3.2 crypto ipsec ike

	crypto ipsec ike {clear sa [peer {any <IPv4 or IPv6 address>} local <IPv4 or IPv6 address>] restart} Manages the IKE (ISAKMP) process or database state.	
Syntax Description	clear	Clears IKE (ISAKMP) peering state
	sa	Clears IKE generated ISAKMP and IPSec security associations (remote peers are affected)
	peer	Clears security associations for the specified IKE peer (remote peers are affected) <ul style="list-style-type: none"> • all—clears security associations for all IKE peerings with a specific local address (remote peers are affected) • IPv4 or IPv6 address—clears security associations for specific IKE peering with a specific local address (remote peers are affected)

	IPv4 or IPv6 address	Clears security associations for the specified IKE peering (remote peer is affected)
	local	Clear security associations for the specified/all IKE peering (remote peer is affected)
	restart	Restarts the IKE (ISAKMP) daemon (clears all IKE state, peers may be affected)
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Example	switch (config)# crypto ipsec ike restart	
Related Commands	show crypto certificate	
Notes		

9.2.3.3 crypto ipsec peer local

	crypto ipsec peer local {enable keying {ike negotiation {ikev1 ikev2} [auth { hmac-sha1 hmac-sha256 hmac-sha512 aes-xcbc} dh-group disable encrypt { 3des-cbc aes-cbc aes-gcm} exchange-mode lifetime local mode peer-identity pfs-group preshared-key prompt-preshared-key transform-set] manual [auth disable encrypt local-spi mode remote-spi]}}	
	Configures IPsec in the system.	
Syntax Description	enable	Enables IPsec peering.
	ike	<p>Configures IPsec peering using IKE ISAKMP to manage SA keys. The following optional parameters are available:</p> <ul style="list-style-type: none"> • auth—configures the authentication algorithm for IPsec peering • dh-group—configures the phase1 Diffie-Hellman group proposed for secure IKE key exchange • disable—configures this IPsec peering administratively disabled • encrypt—configures the encryption algorithm for IPsec peering • exchange-mode—configures the IKE key exchange mode to propose for peering • lifetime—configures the SA lifetime to propose for this IPsec peering • local-identity—configures the ISAKMP payload identification value to send as local endpoint's identity • mode—configures the peering mode for this IPsec peering • peer-identity—configures the identification value to match against the peer's ISAKMP payload identification • pfs-group—configures the phase2 PFS (Perfect Forward Secrecy) group to propose for Diffie-Hellman exchange for this IPsec peering • preshared-key—configures the IKE pre-shared key for the IPsec peering • prompt-preshared-key—prompts for the pre-shared key, rather than entering it on the command line • transform-set—configures transform proposal parameters

	keying	Configures key management for this IPsec peering. <ul style="list-style-type: none"> • auth—configures the authentication algorithm for this IPsec peering • disable—configures this IPsec peering administratively disabled • encrypt—configures the encryption algorithm for this IPsec peering • local-spi—configures the local SPI for this manual IPsec peering • mode—configures the peering mode for this IPsec peering • remote-spi—configures the remote SPI for this manual IPsec peering
	manual	Configures IPsec peering using manual keys.
Default	N/A	
Configuration Mode	config	
History	3.2.3000 3.9.3100: Added support for IKEv2 and new ciphers	
Example	switch (config)# crypto ipsec peer 10.10.10.10 local 10.7.34.139 enable	
Related Commands	show crypto certificate	
Notes	As of version 3.9.3100, NULL will not be supported as an authentication or encryption algorithm for IPsec peering. New ciphers are supported (hmac-sha512 and aes-xcbc for authentication and aes-gcm for encryption. 1, 2, 5, 22, 23, 24 pfs/dh-groups will not be supported, while 19, 20, 21 will be supported only with IKEv2. The transform-set options ah-and-esp-ah are no longer supported. Libreswan is used instead of openswan.	

9.2.3.4 crypto certificate ca-list

	crypto certificate ca-list [default-ca-list name {<cert-name> system-self-signed}] no crypto certificate ca-list [default-ca-list name {<cert-name> system-self-signed}] Adds the specified CA certificate to the default CA certificate list. The no form of the command removes the certificate from the default CA certificate list.	
Syntax Description	cert-name	The name of the certificate
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Example	switch (config) # crypto certificate default-cert name test	
Related Commands	show crypto certificate	
Notes	<ul style="list-style-type: none"> • 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 	

9.2.3.5 crypto certificate default-cert

	crypto certificate default-cert name {<cert-name> system-self-signed} no crypto certificate default-cert name {<cert-name> system-self-signed} Designates the named certificate as the global default certificate role for authentication of this system to clients. The no form of the command reverts the default-cert name to “system-self-signed” (the “cert-name” value is optional and ignored).	
Syntax Description	cert-name	The name of the certificate
Default	N/A	
Configuration Mode	config	
History	3.2.3000	
Example	switch (config) # crypto certificate default-cert name test	
Related Commands	show crypto certificate	
Notes	<ul style="list-style-type: none"> • 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 	

9.2.3.6 crypto certificate generation

	crypto certificate generation default {country-code days-valid > ca-valid <true/false> email-addr hash-algorithm {sha1 sha256} key-size-bits locality org-unit organization state-or-prov} Configures default values for certificate generation.	
Syntax Description	country-code	Configures the default certificate value for country code with a two-alphanumeric-character code or -- for none.
	days-valid	Configures the default certificate valid days Default value: 365 days
	email-addr	Configures the default certificate value for email address
	hash-algorithm {sha1 sha256}	Configures the default certificate hashing algorithm
	key-size-bits	Configures the default certificate value for private key size (private key length in bits—at least 1024, but 2048 is strongly recommended)
	locality	Configures the default certificate value for locality
	org-unit	Configures the default certificate value for organizational unit
	organization	Configures the default certificate value for the organization name
	state-or-prov	Configures the default certificate value for state or province
	ca-valid {true false}	Configures the default certificate CA Basic Constraints flag set to TRUE/FALSE
Default	hash-algorithm - sha1	
Configuration Mode	config	

History	3.2.1000 3.3.4350: Added "hash-algorithm" parameter 3.6.4000: Added "days-valid" parameter 3.8.2100: Added "ca-valid" parameter
Example	<code>switch (config) # crypto certificate generation default hash-algorithm sha256</code>
Related Commands	<code>show crypto certificate</code>
Notes	

9.2.3.7 crypto certificate name

	<p><code>crypto certificate name {<cert-name> system-self-signed} {comment <new comment> generate selfsigned [comment <cert-comment> common-name <domain> country-code <code> days-valid <days> ca-valid <true/false> email-addr <address> hash-algorithm {sha1 sha256} key-size-bits <bits> locality <name> org-unit <name> organization <name> serial-num <number> state-or-prov <name>] private-key pem <PEM string> prompt-private-key public-cert [comment <comment string> pem <PEM string>] regenerate days-valid <days> ca-valid <true/false> rename <new name>}</code></p> <p><code>no crypto certificate name <cert-name></code></p> <p>Configures default values for certificate generation. The no form of the command clears/deletes certain certificate settings.</p>	
Syntax Description	<code>cert-name</code>	Unique name by which the certificate is identified.
	<code>comment</code>	Specifies a certificate comment.
	<code>generate self-signed</code>	<p>Generates certificates. This option has the following parameters which may be entered sequentially in any order:</p> <ul style="list-style-type: none"> <code>comment</code>—specifies a certificate comment (free string) <code>common-name</code>—specifies the common name of the issuer and subject (e.g. a domain name) <code>country-code</code>—specifies the country codwo-alphanumeric-character country code, or "--" for none) <code>days-valid</code>—specifies the number of days the certificate is valid <code>email-addr</code>—specifies the email address <code>hash-algorithm</code>—specifies the hashing function used for signature algorithm. Default value is SHA256. <code>key-size-bits</code>—specifies the size of the private key in bits (private key length in bits - at least 1024 but 2048 is strongly recommended) <code>locality</code>—specifies the locality name <code>org-unit</code>—specifies the organizational unit name <code>organization</code>—specifies the organization name <code>serial-num</code>—specifies the serial number for the certificate (a lower-case hexadecimal serial number prefixed with "0x") <code>state-or-prov</code>—specifies the state or province name <code>ca-valid</code>—Specifies certificate CA Basic Constraints flag set to TRUE/FALSE
	<code>private-key pem</code>	Specifies certificate contents in PEM format
	<code>prompt-private-key</code>	Prompts for certificate private key with secure echo
	<code>public-cert</code>	Installs a certificate
	<code>regenerate</code>	Regenerates the named certificate using configured certificate generation default values for the specified validity period

	rename	Renames the certificate
Default	N/A	
Configuration Mode	config	
History	3.2.3000 3.3.4402: Added "hash-algorithm" parameter 3.6.4000: Added "days-valid" parameter 3.8.2100: Added "ca-valid" parameter	
Example	switch (config) # crypto certificate name system-self-signed generate self-signed hash-algorithm sha256	
Related Commands	show crypto certificate	
Notes		

9.2.3.8 crypto certificate system-self-signed

	crypto certificate system-self-signed regenerate {[days-valid <days>] ca-valid <true/false>} Configures default values for certificate generation.	
Syntax Description	days-valid	Specifies the number of days the certificate is valid
	ca-valid	Specifies certificate CA Basic Constraints flag set to TRUE/FALSE
Default	N/A	
Configuration Mode	config	
History	3.2.1000 3.8.2100: Added the ca-valid option	
Example	switch (config) # crypto certificate system-self-signed regenerate days-valid 3 switch (config) # crypto certificate system-self-signed regenerate ca-valid false	
Related Commands	show crypto certificate	
Notes		

9.2.3.9 show crypto certificate

	show crypto certificate [detail public-pem default-cert [detail public-pem] [name <cert-name> [detail public-pem] ca-list [default-ca-list]] Displays information about all certificates in the certificate database.	
Syntax Description	ca-list	Displays the list of supplemental certificates configured for the global default system CA certificate role
	default-ca-list	Displays information about the currently configured default certificates of the CA list
	default-cert	Displays information about the currently configured default certificate
	detail	Displays all attributes related to the certificate
	name	Displays information about the certificate specified
	public-pem	Displays the uninterpreted public certificate as a PEM formatted data string

Default	N/A
Configuration Mode	config
History	3.2.1000 3.8.2100: Updated output
Example	
<pre>switch (config) # show crypto certificate Certificate with name 'system-self-signed' (default-cert) Comment: system-generated self-signed certificate Private Key: present Serial Number: 0x546c935511bcafc21ac0e8249fbe0844 SHA-1 Fingerprint: fe6df38dd26801971cb2d44f62dbe492b6063c5f Validity: Starts: 2012/12/02 13:45:05 Expires: 2013/12/02 13:45:05 Subject: Common Name: IBM-DEV-Bay4 Country: IS State or Province: Locality: Organization: Organizational Unit: E-mail Address: Issuer: Common Name: IBM-DEV-Bay4 Country: IS State or Province: Locality: Organization: Organizational Unit: E-mail Address: X509 Extensions: Basic Constraints: CA: TRUE</pre>	
Related Commands	
Notes	

9.2.3.10 show crypto encrypt-data

	show encrypt-data Displays sensitive data encryption information.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.1002
Example	<pre>switch (config)# show crypto encrypt-data Sensitive files encryption: Status: enabled Key location: usb Cipher: aes256</pre>
Related Commands	
Notes	

9.2.3.11 show crypto ipsec

	<code>show crypto ipsec [brief configured ike policy sa]</code> Displays information ipsec configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.2.1000
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.</pre>
Related Commands	
Notes	

10 InfiniBand Switching

The following pages provide information on configuring InfiniBand protocols and features.

- [Node Name](#)
- [Fabric](#)
- [IB Router](#)
- [InfiniBand Interface](#)
- [Subnet Manager](#)
- [Subnet Manager High Availability](#)

10.1 Node Name

10.1.1 Node Name Commands

- [10.1.1 Node Name Commands](#)
 - [10.1.1.1 ib nodename](#)
 - [10.1.1.2 show ib nodename](#)

10.1.1.1 ib nodename

	ib nodename <guid> name <name> no ib nodename <guid> Maps GUID and node name. The no form of the command unmaps the GUID and node name.	
Syntax Description	guid	System GUID
	name	User defined string
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib nodename 00:00:00:00:60:04:03:30 name my-name	
Related Commands		
Notes	If an entry with the same GUID exists, the existing name will be replaced with a new name.	

10.1.1.2 show ib nodename

	show ib nodename Displays nodename and GUID information.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000

Example	switch (config) # show ib nodename GUID='00:00:00:00:60:04:03:30', name='my-name', discovered='no'
Related Commands	ib nodename
Notes	

10.2 Fabric



10.2.1 Fabric Commands

- [10.2.1 Fabric Commands](#)
 - [10.2.1.1 show guides](#)
 - [10.2.1.2 show system guid](#)
 - [10.2.1.3 show lids](#)

10.2.1.1 show guides

	show guides Displays GUIDs per ASIC in the chassis.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.4.2008	Updated example
	3.6.1002	Updated example
Example	<pre>switch (config) # show guides ===== Module Device IB Subnet GUID ===== SYSTEM - - E4:1D:2D:03:00:2E:49:40 MGMT SIB infiniband-default E4:1D:2D:03:00:2E:49:40 MGMT SIB infiniband-1 E4:1D:2D:03:00:2E:49:41 MGMT SIB infiniband-2 E4:1D:2D:03:00:2E:49:42</pre>	
Related Commands		
Notes		

10.2.1.2 show system guid

	show {guids system guid} Displays the system GUID.
Syntax Description	N/A
Default	N/A
Configuration Mode	config

History	3.1.0000
Example	switch (config) # show system guid 00:02:C9:03:00:43:D9:00
Related Commands	
Notes	

10.2.1.3 show lids

	show lids Displays the LIDs of each module in the switch system.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.4.2008	Updated example
	3.6.1002	Updated example
Example	<pre>switch (config) # show lids ===== Module Device IB Subnet LID ===== MGMT SIB infiniband-default 1 MGMT SIB infiniband-1 8 MGMT SIB infiniband-2 3</pre>	
Related Commands		
Notes		

10.3 IB Router



IB router provides the ability to send traffic between two or more IB subnets thereby potentially expanding the size of the network to over 40k end-ports, enabling separation and fault resilience between islands and IB subnets, and enabling connection to different topologies used by different subnets.


The forwarding between the InfiniBand subnets is performed using GRH (global route header) lookup.

IB router capabilities are supported only on QM9700 switch systems.

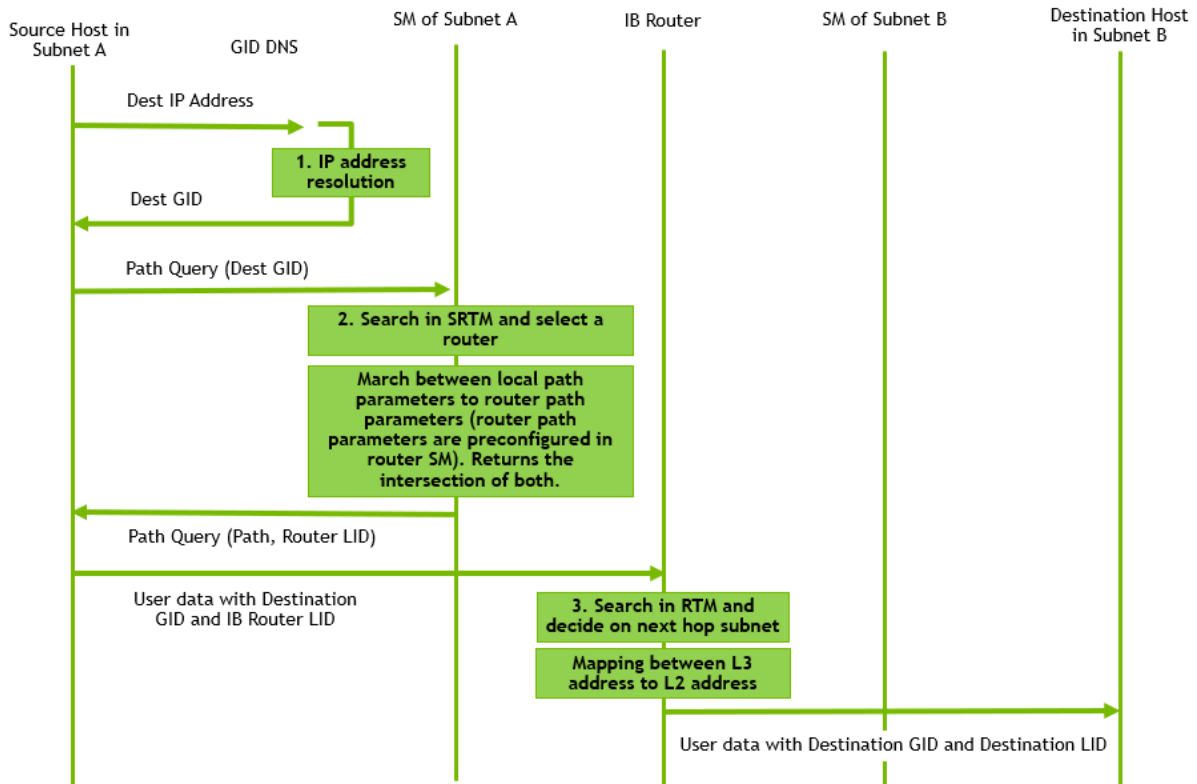
The IB router's basic functionality includes:

- Removal of current L2 LRH (local routing header)
- Routing table lookup—using GID from GRH
- Building new LRH according to the destination and the routing table

The DLID in the new LRH is built using simplified GID-to-LID mapping (where LID = 16 LSB bits of GID) thereby not requiring to send for ARP query/lookup.



For this to work, the SM allocates an alias GID for each host in the fabric where the alias GID = {subnet prefix[127:64], reserved[63:16], LID[15:0]}. Hosts should use alias GIDs in order to transmit traffic to peers on remote subnets.



For more information on IB router architecture and functionality, please refer to the community post [IB Router Architecture and Functionality](#).

IB router requires HCA configuration such as SM, partition key, MPI, GID translation, and more. To learn more about these configurations, please refer to the following community posts:

- [INFINIBAND FABRIC CONFIGURATION WITH QUANTUM2 INFINIBAND ROUTER](#)
- [INFINIBAND FABRIC CONFIGURATION WITH QUANTUM2 INFINIBAND ROUTER—PART 2](#)

The minimal UFM SM version for NDR multi-SWID functionality is 5.15 and above.

10.3.1 Configuring IB Router

10.3.1.1 Prerequisites

1. Check system capabilities to make sure IB L3 is supported. Run:

```
switch (config) # show system capabilities
IB: Supported, L2, L3, Adaptive Routing, Split Ready
Max SM nodes: 2048
IB Max licensed speed: NDR
```

Please notice the second line in the output.

2. Configure system profile to multi-switch with 2 SWIDs.

```
switch (config) # system profile ib num-of-swids 2 ib-router
```

Note that all interfaces are mapped to subnet infiniband-default.

3. Verify system profile configuration.

```
switch (config) # show system profile
Profile           : ib
Number of SWIDs   : 2
Adaptive Routing  : yes
Adaptive Routing Groups: 256
Split Ready       : no
IB Routing        : yes
```

Note the number of SWIDs configured and that IB Routing is set to “yes”.

10.3.1.2 Configuring IB Router

1. Map an interface to a SWID.

```
switch (config) # interface ib 1/1/1 switchport access subnet infiniband-default force
switch (config) # interface ib 1/1/2 switchport access subnet infiniband-1 force
```

2. Verify SWID configuration on the aforementioned interfaces.

```
switch (config) # show interfaces ib status
-----
Interface      Description      IB Subnet      Speed      Current line rate  Logical port state
Physical port state
-----
IB1/1/1
LinkUp
IB1/1/2
LinkUp
IB1/2/1
Polling
IB1/2/2
Polling
IB1/3/1
Polling
IB1/3/2
Polling
-----
```

Interface	Description	IB Subnet	Speed	Current line rate	Logical port state
IB1/1/1		infiniband-default	ndr	400.0 Gbps	Initialize
IB1/1/2		infiniband-1	ndr	400.0 Gbps	Initialize
IB1/2/1		infiniband-default	-	-	Down
IB1/2/2		infiniband-default	-	-	Down
IB1/3/1		infiniband-default	-	-	Down
IB1/3/2		infiniband-default	-	-	Down

```

IB1/4/1          infiniband-default - - - Down
Polling
IB1/4/2          infiniband-default - - - Down
Polling
IB1/5/1          infiniband-default - - - Down
Polling
..

```

3. Configure and enable InfiniBand router.

```

switch (config) # ib router
switch (config) # no ib router shutdown

```

4. Enable IB subnet interface.

```

switch (config) # no interface ib-subnet infiniband-default shutdown
switch (config) # no interface ib-subnet infiniband-1 shutdown

```

5. Verify configuration.

```

switch (config) # show ib router
Routing state: enabled
-----
IB subnet          Routing enabled
-----
infiniband-default  enabled
infiniband-1       enabled
switch (config) # show interfaces ib-subnet infiniband-default
infiniband-default state:
  GUID          : 90:0A:84:03:00:40:C9:C8
  Alias GUID    : N/A
  LID           : 4
  Subnet prefix : FE:C0:00:00:00:00:11
  Physical state : LinkUp
  Logical state  : Active
  L3 interface state: Up

```

10.3.1.3 IP to GUID Resolution

1. Go to the following Github: <https://github.com/Mellanox/ip2gid>
2. Clone the Git repository
3. Compile and run on each node in the fabric
4. Change the device MAC address of the IPoIB device to be based on the alias GUID and not the GUID.

For example, # echo fec0:0000:0000:0003:0014:0500:0000:0001 > /sys/class/net/ib0/
set_mac

where fe:c0:00:00:00:00:02:00:14:05:00:00:00:00:01 is the alias GUID given by the SM to that node.

5. Add route using "ip route add" command to the relevant hosts.


```

# ifconfig ib0 12.0.3.1/24 --> set ip for ib0
# ip route add 12.0.1.0/24 via 12.3.0.250 --> adding route to hosts with 12.1.xxx.xxx IP
# ip route add 12.0.2.0/24 via 12.3.0.250 --> adding route to hosts with 12.2.xxx.xxx IP

```

10.3.2 Subnet Prefix Checking

Subnet prefix checking only applies for when MLNX OS subnet manager is running in the InfiniBand fabric.

Subnet manager can't be started on the switch with IB router functionality enabled.

To allow InfiniBand routing, the subnet prefixes in all routable subnets must be in site-local format - fe:c0:00:00:00:00:xx:xx (e.g. fe:c0:00:00:00:00:00:01).

By default, the command which defines the subnet prefix of the Infiniband subnet, validates the subnet prefix before allowing the change.

For proper IB management of the Infiniband fabric including IB routers, the recommended order of commands is as follows:

- `ib sm subnet-prefix` - configures the subnet prefix
- `ib sm rtr-aguid-enable <1 | 2>` - enables support for Host alias GIDs needed for sending routable traffic.
- `ib sm enable` - start SM on this node or any node in cluster.

To disable subnet prefix checking

1. Verify the status of subnet prefix override. Run:

```
switch (config) # show ib sm subnet-prefix-override
enable
```

2. If enabled, disable subnet-prefix-override. Run:

```
switch (config) # ib sm subnet-prefix-override
```

3. Verify configuration. Run:

```
switch (config) # show ib sm subnet-prefix-override
disable
```

10.3.3 IB Router Commands

- [10.3.1 Configuring IB Router](#)
 - [10.3.1.1 Prerequisites](#)
 - [10.3.1.2 Configuring IB Router](#)
 - [10.3.1.3 IP to GID Resolution](#)
- [10.3.2 Subnet Prefix Checking](#)
- [10.3.3 IB Router Commands](#)
 - [10.3.3.1 ib router](#)
 - [10.3.3.2 ib router shutdown](#)
 - [10.3.3.3 interface ib-subnet](#)
 - [10.3.3.4 interface ib-subnet shutdown](#)
 - [10.3.3.5 show ib router](#)
 - [10.3.3.6 show interfaces ib-subnet](#)

10.3.3.1 ib router

	<code>ib router</code> <code>no ib router</code> Enables the set of commands that allow control of IB router functionality. The no form of the command disables IB router commands and removes all related configurations.
--	--

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.0500
Example	<code>switch (config) # ib router</code>
Related Commands	system profile
Notes	

10.3.3.2 ib router shutdown

	<code>ib router shutdown</code> <code>no ib router shutdown</code> Disables IB router. The no form of the command enables IB router.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config	
History	3.6.0500	
Example	<code>switch (config) # no ib router shutdown</code>	
Related Commands		
Notes	This command does not clear IB router configuration	

10.3.3.3 interface ib-subnet

	<code>interface ib-subnet <swid-name></code> <code>no interface ib-subnet <swid-name></code> Creates routing on IB router subnet. The no form of the command removes routing on router interface.	
Syntax Description	swid-name	Name of the SWID: infiniband-default, infiniband-1...infiniband-5
Default	N/A	
Configuration Mode	config	
History	3.6.0500	
Example	<code>switch (config) # interface ib-subnet infiniband-3</code>	
Related Commands	system profile	
Notes	The maximum number of SWIDs depends on the number of SWIDs defined in the profile	

10.3.3.4 interface ib-subnet shutdown

	interface ib-subnet <swid-name> shutdown no interface ib-subnet <swid-name> shutdown Disables routing on IB router subnet. The no form of the command enables routing on router interface.	
Syntax Description	swid-name	Name of the SWID: infiniband-default, infiniband-1...infiniband-5
	shutdown	Admin down on router interface Admin up on router interface with no form of command
Default	Disabled	
Configuration Mode	config	
History	3.6.0500	
Example	<pre>switch (config) # no interface ib-subnet infiniband-3 shutdown</pre>	
Related Commands		
Notes		

10.3.3.5 show ib router

	show ib router Displays current IB router functionality.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.0500	
Example	<pre>switch (config) # show ib router Routing state: enabled IB Subnet Routing infiniband-default enabled infiniband-1 disabled infiniband-2 enabled infiniband-3 enabled</pre>	
Related Commands		
Notes		

10.3.3.6 show interfaces ib-subnet

	show interfaces ib-subnet [<swid-name>] [brief] Displays statistics of one or all IB subnets with enabled IB routing.	
Syntax Description	swid-name	Name of the SWID: infiniband-default, infiniband-1...infiniband-5
	brief	Displays output in a table format
Default	Disabled	
Configuration Mode	config	

History	3.6.0500
Example	<pre>switch (config) # show interfaces ib-subnet infiniband-3 infiniband-3 state: GUID : F4:52:14:03:00:6E:F2:8B Alias GUID : N/A LID : 10 Subnet prefix : FE:C0:00:00:00:00:08 Physical state : LinkUp Logical state : Active L3 interface state : Up</pre>
Related Commands	
Notes	

10.4 InfiniBand Interface



10.4.1 Transceiver Information

MLNX-OS offers the option of viewing the transceiver information of a module or cable connected to a specific interface. The information is a set of read-only parameters burned onto the EEPROM of the transceiver by the manufacture. The parameters include identifier (connector type), cable type, speed and additional inventory attributes.

To display transceiver information of a specific interface, run:

```
switch (config) # show interfaces ib 1/36 transceiver
Slot 1 port 36 state
  identifier          : QSFP+
  cable/module type   : Passive copper, unequalized
  infiniband speeds    : SDR , DDR , QDR , FDR , HDR , NDR
  vendor              : Mellanox
  cable length        : 2m
  part number         : MC2207130-0A1
  revision            : A3
  serial number       : MT1324VS02215
```

The indicated cable length is rounded up to the nearest natural number.

10.4.2 High Power Transceivers

NVIDIA switch systems offer high power transceiver (e.g. LR4) support on all ports of the Switch-IB® family switch systems.

If a high power transceiver (e.g. LR4) is inserted to a port that does not support it, the link does not go up, and the following warning message is displayed: “Warning: High power transceiver is not supported” when the command “show interfaces ib” is run.

10.4.3 Forward Error Correction

Forward Error Correction (FEC) mechanism adds extra data to the transmitted information. The receiving device uses this additional data to verify that the received data contains no errors. If the

receiving side discovers errors within the received data it is able to correct some of these errors. The number of errors that can be corrected depends on the FEC algorithm.

Switch-IB® EDR (100Gb/s) NVIDIA-to-NVIDIA InfiniBand connections enable standard low-latency Reed Solomon (LL RS) FEC on active optical cables longer than 30 meters and passive copper cables longer than 2m.

10.4.4 Breakout Cables

This feature is available only for Quantum and Quantum-2 based switch systems (including modular systems).

To split a port in an externally managed 1U switch system, please refer to section “Using mlxconfig to Split a Port in a Remotely Managed Switch” in the MFT documentation.

In-band management is blocked by default on MLNX-OS switches, but can be enabled via the CLI command `system manage inband-ib`. Once enabled, interaction between the different sources of commands is possible. That is, you may split a port via MAD and display the results on CLI, or split a port via CLI and query the result via MAD. Logically, it is similar to having two CLI users.

The break-out cable is a unique NVIDIA capability, where a single physical quad-lane QSFP port is divided into 2 dual-lane ports. It maximizes the flexibility of the end user to use the NVIDIA switch with a combination of dual-lane and quad-lane interfaces according to the specific requirements of its network. All system ports may be split into 2-lane ports and, on modular systems, only external ports may be split. Splitting a port changes the notation of that port from x/y to $x/y/z$ with “ x/y ” indicating the previous notation of the port prior to the split and “ z ” indicating the number of the resulting single-lane port (1,2). Each sub-physical port is then handled as an individual port. For example, splitting port 5 into 2 lanes gives the following new ports: 1/5/1 & 1/5/2 and on modular systems, splitting port 5 of device 1 on leaf 1 will give the following ports: 1/1/5/1 and 1/1/5/2.

For Quantum-2 based systems, there will be an additional hierarchy—the cage level, which, for Quantum-2 systems, contains 2 physical ports.

As such, a representation of an interface for Quantum-2 systems will be `<Asic/Cage/Port>` (e.g., “interface ib 1/3/1 and interface ib 1/3/2” represents 2 ports 1 & 2 that are located within cage 3 of the system).

For split interfaces on Quantum-2 systems, the representation will be `<Asic/Cage/Port/split >` (e.g., “interface ib 1/3/1/1 and interface ib 1/3/1/2” represents split ports 1 & 2 that are split of port 1 located within cage 3 of the system).

For example, QM9700 system has 32 cages, but 64 ports are represented from “interface ib 1/1/1-1/32/2”.

Splitting the interface deletes all configuration on that interface.

In order to use this feature, the system’s profile must be configured to “ib split-ready” as described in section [“Changing System Profile to Allow for Split-Ready Configuration”](#) using the command `system profile`.

On modular systems, both managements should be configured with split-ready profile.

10.4.4.1 Changing System Profile to Allow for Split-Ready Configuration

If system does not have split-ready configuration, change its profile to allow for it:

1. Change the system's profile to "ib split-ready". Run:

```
switch (config) # system profile ib split-ready
Warning - confirming will cause system reboot and all configuration will be deleted
Type 'yes' to confirm profile change: yes
```

On modular systems, system's profile need to be changed to "ib split-ready" for both managements simultaneously.

2. Verify system profile configuration. Run:

```
switch (config) # show system profile
Profile: ib
Number of SWIDs: 1
Adaptive Routing: yes
Adaptive Routing Groups: 1792
Split Ready: yes
IB Routing: no
```

10.4.4.2 Changing the Module Type to a Split Mode

To split an interface:

1. Shut down the interface.

Examples:

```
switch (config)# interface ib 1/4 shutdown

or

switch (config) # interface ib 1/4
switch (config interface ib 1/4) # shutdown
```

2. Split the ports as desired. Run:

```
switch (config interface ib 1/4) # module-type qsfpsplit-2
```

3. New ports can be shown by the interfaces IB status command:

```
switch (config) # show interfaces ib status
Interface  Description  IB Subnet  Speed  Current line rate  Logical port state
Physical port state
-----
IB1/1/1    infiniband-default  edr      25.0 Gbps  Active
LinkUp
IB1/1/2    infiniband-default  -        -        Down
Polling
IB1/2      infiniband-default  edr      100.0 Gbps  Active
LinkUp
IB1/3      infiniband-default  -        -        Down
Polling
IB1/4/1    infiniband-default  -        -        Down
Polling
IB1/4/2    infiniband-default  -        -        Down
Polling
IB1/5      infiniband-default  -        -        Down
Polling
```

IB1/6 Polling	infiniband-default	-	-	Down
IB1/7 Polling	infiniband-default	-	-	Down
IB1/8 Polling	infiniband-default	-	-	Down

The above examples were executed on 1U systems, but are available also for modular systems.

10.4.4.3 Unsplitting a Split Port

To unsplit a split port:

1. Shut down all of the split ports. Run:

```
switch (config) # interface ib 1/4/1
switch (config interface ib 1/4/1) # shutdown
switch (config interface ib 1/4/1) # exit
switch (config) # interface ib 1/4/2
switch (config interface ib 1/4/2) # shutdown
switch (config interface ib 1/4/2) # exit
```

2. From the first member of the split (1/4/1), change the module-type back to QSFP. Run:

```
switch (config interface ib 1/4/1) # module-type qsfp
```

For both split and unsplit operations, a “force” option may be utilized to allow the operation without previously shutting down the ports.

The above examples were executed on 1U systems, but are available also for modular systems.

10.4.5 InfiniBand Interface Commands



- [10.4.5.1 interface ib](#)
- [10.4.5.2 module-type](#)
- [10.4.5.3 interface ib port-type split-2](#)
- [10.4.5.4 mtu](#)
- [10.4.5.5 shutdown](#)
- [10.4.5.6 description](#)
- [10.4.5.7 speed](#)
- [10.4.5.8 op-vls](#)
- [10.4.5.9 width](#)
- [10.4.5.10 clear counters](#)
- [10.4.5.11 interface ib internal notification link-speed-mismatch](#)
- [10.4.5.12 interfaces ib internal notification link-state-change](#)
- [10.4.5.13 switchport access subnet](#)

- [10.4.5.14 show interfaces ib](#)
- [10.4.5.15 show interfaces ib status](#)
- [10.4.5.16 show interfaces ib internal](#)
- [10.4.5.17 show interfaces ib internal capabilities](#)
- [10.4.5.18 show interfaces ib internal llr](#)
- [10.4.5.19 show interfaces ib internal status](#)
- [10.4.5.20 show interfaces ib transceiver](#)
- [10.4.5.21 show interfaces ib transceiver diagnostics](#)
- [10.4.5.22 show interfaces ib transceiver raw](#)

10.4.5.1 interface ib

	interface ib [internal] {<inf> <inf-range>} Enters the InfiniBand interface configuration mode.	
Syntax Description	[internal] <inf>	For 1U switches: interface 1/<interface> For modular switches: <ul style="list-style-type: none"> • interface ib <interface> • interface ib internal leaf <interface> • interface ib internal spine <interface>
	inf-range	Enters the configuration mode of a range of interfaces Format: <slot>/<port>[-<slot>/<port>]
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.4.2008	Added internal leaf and spine options
Example	switch (config) # interface ib 1/1 switch (config interface ib 1/1) #	
Related Commands	show interfaces ib	
Notes	Interface range (inf-range) option is not valid on modular switch systems	

10.4.5.2 module-type

	module-type <qsf type> Split or unsplit the interface.	
Syntax Description	qsf	Unsplits the interface
	qsf-split-2	Splits the interface
Default	Disabled	
Configuration Mode	config interface ib	
History	3.8.1000	Added splitting capability on 1U
	3.9.0300	Added splitting capability on modular systems
Example	switch (config)# interface ib 1/1module-type qsf-split-2switch (config) #interface ib 1/1 module-type qsf	

Related Commands	show interfaces ib
Notes	Available only for systems configured with split-ready profile. See " Break-Out Cables " section for more information.

10.4.5.3 interface ib port-type split-2

	interface ib <port> port-type split-2 [force] no interface ib <port> port-type [force] Splits selected port to 2. The no form of the command unsplit the selected port.	
Syntax Description	port	Selected port.
	force	In force mode, the port does not need to be disabled before split. If force parameter is not set, the ports needs to be disabled.
Default	Disabled	
Configuration Mode	config interface ib	
History	3.10.0100	
Example	<pre>switch (config) # interface ib 1/8/2 port-type split-2 force switch (config) # no interface ib 1/6/2 port-type</pre>	
Related Commands		
Notes	Available only when split profile is configured.	

10.4.5.4 mtu

	mtu <frame-size> no interface ib mtu Configures the Maximum Transmission Unit (MTU) frame size for the interface. The no form of the command returns the field value to its default.	
Syntax Description	frame-size	MTU size in bytes Possible values: 256, 512 ,1K, 2K, 4K (K =1024)
Default	4096 bytes	
Configuration Mode	config interface ib	
History	3.1.0000	
	3.9.3100	Added the no form of the command
Example	<pre>switch (config interface ib 1/1) # mtu 4K</pre>	
Related Commands	show interfaces ib	
Notes		

10.4.5.5 shutdown

	shutdown no shutdown Disables the interface. The no form of the command enables the interface.	
--	---	--

Syntax Description	N/A
Default	Enabled
Configuration Mode	config interface ib
History	3.1.0000
Example	switch (config interface ib 1/1) # shutdown
Related Commands	show interfaces ib
Notes	

10.4.5.6 description

	description <string> no description Configures an interface description. The no form of the command deletes interface description.	
Syntax Description	string	40 bytes
Default	""	
Configuration Mode	config interface ib	
History	3.1.0000	
Example	switch (config interface ib 1/1) # description my_interface	
Related Commands	show interfaces ib	
Notes		

10.4.5.7 speed

	speed <port speed> [force] Configures the speed negotiation of the interface.	
Syntax Description	port speed	The following options are available: <ul style="list-style-type: none"> • sdr—10.0Gb/s rate on 4 lane width • ddr—20.0Gb/s rate on 4 lane width • qdr—40.0Gb/s rate on 4 lane width • fdr10—40.0Gb/s rate on 4 lane width • fdr—56.0Gb/s rate on 4 lane width • edr—100.0Gb/s rate on 4 lane width • hdr—200.0 Gb/s rate on 4 lane width • ndr—400.0 Gb/s rate on 4 lane width
	force	Forces configuration of speed-list not containing SDR bit
Default	Depends on the port module type, not all interfaces support all speed options	
Configuration Mode	config interface ib	
History	3.1.0000	
	3.4.1604	Updated Syntax Description and Example
	3.8.2000	Updated port speed in Syntax Description and Example
	3.10.0100	Added NDR speed

Example	<code>switch (config interface ib 1/1) # speed fdr edr hdr ndr</code>
Related Commands	<code>show interfaces ib</code>
Notes	<ul style="list-style-type: none"> • This command is backwards compatible so old configuration file containing this command with the old form (with legal bit mask) are still supported • Configuring more than one speed is possible by typing in consecutive speed names separated by spaces • If the speed-options list does not include SDR speed, it is configured automatically. However, if the force option is used (supported on FDR10 only), SDR is not configured. • If the other side of the link is a ConnectX®-3 device, to allow the link to raise in FDR speed, QDR speed must also be allowed • Force parameter is only allowed for SIB2 systems for EDR speed option.

10.4.5.8 op-vls

	<code>op-vls <value></code> <code>no op-vls</code> Configures the operational VLs of the interface. The no form of the command sets the operational VLs to its default value.	
Syntax Description	value	Possible value for operational VLs <ul style="list-style-type: none"> • 1 VL0 • 2 VL0, VL1 • 4 VL0-VL3 • 8 VL0-VL7
Default	8 (VL0 - VL7)	
Configuration Mode	<code>config interface ib</code>	
History	3.1.0000	
Example	<code>witch (config interface ib 1/1) # op-vls 1</code>	
Related Commands	<code>how interfaces ib</code>	
Notes		

10.4.5.9 width

	<code>width <value></code> <code>no width</code> Sets the width of the interface. The no form of the command resets the parameter to its default value.	
Syntax Description	value	Possible value for width for an unsplit port: <ul style="list-style-type: none"> • 1-1X • 3-1X, 2X • 5-1X, 4X • 7-1X, 2X, 4X
Default	7	
Configuration Mode	<code>config interface ib</code>	
History	3.1.0000	
Example	<code>switch (config interface ib 1/1) # width 1</code>	
Related Commands	<code>show interfaces ib</code>	

Notes	
-------	--

10.4.5.10 clear counters

	clear counters Clears the interface counters.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config interface ib	
History	3.1.0000	
Example	switch (config interface ib 1/1) # clear counters	
Related Commands	show interfaces ib	
Notes		

10.4.5.11 interface ib internal notification link-speed-mismatch

	interface ib internal notification link-speed-mismatch [<time>] no interface ib internal notification link-speed-mismatch Enables notifications on internal link speed mismatch in SNMP. The no form of the command disables notifications on internal links speed mismatch in SNMP.	
Syntax Description	time	In hours. Enables periodic notifications (traps and log) on internal link speed mismatch status. "0" disables the feature.
Default	Disabled	
Configuration Mode	config	
History	3.4.3000	
Example	switch (config) # interface ib internal link-speed-mismatch 6	
Related Commands	show interfaces ib internal notification	
Notes	Link-speed-mismatch shows internal link entries in the iVPITable	

10.4.5.12 interfaces ib internal notification link-state-change

	interfaces ib internal notification link-state-change no interfaces ib internal notification link-state-change Enables notifications on internal links state change in SNMP. The no form of the command disables notifications on internal links state change in SNMP.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config	
History	3.3.4318	

Example	<code>switch (config) # interfaces ib internal notification</code>
Related Commands	<code>show interfaces ib internal notification</code>
Notes	Link-state-change shows internal link entries in the ifTable and the ifXTable

10.4.5.13 switchport access subnet

	<code>switchport access subnet <swid-name> [force]</code> <code>no switchport access subnet <swid-name> [force]</code> Maps interface to SWID. The no form of the command unmaps an interface from a SWID.	
Syntax Description	swid-name	Name of the SWID: infiniband-default, infiniband-1...infiniband-5
	force	Applies configuration without the need to shutdown the interface before running command
Default	Unmapped	
Configuration Mode	config interface ib	
History	3.6.0500	
Example	<code>switch (config interface ib1/36) # switchport access subnet infiniband-1</code>	
Related Commands		
Notes	<ul style="list-style-type: none"> • Mapping an interface automatically enables it • Remapping an interface resets all its configuration except for interface description • Unmapping an interface resets all its configuration except for interface description • An interface needs to be disabled before remapping/unmapping unless the “force” parameter is used 	

10.4.5.14 show interfaces ib

	<code>show interfaces ib <inf></code> Displays the configuration and status for the interface.	
Syntax Description	internal	Internal interfaces
	inf	<ul style="list-style-type: none"> • Slot/Port (i.e. 1/1) • LXX/SXX (i.1 L01 or S01)
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.4.1604	Updated example
	3.6.1002	Updated example
	3.6.6105	Updated example
	3.9.1300	Updated output
	3.9.3100	Updated field 'VL capabilities' to 'VL admin capabilities'

Example	<pre> switch (config) # show interfaces ib 1/1 IB1/1 state: Logical port state : Down Physical port state : Polling Current line rate : - Supported speeds : sdr, ddr, qdr, fdr10, fdr, edr Speed : - Supported widths : 1X, 4X Width : 4X Max supported MTUs : 4096 MTU : 0 VL admin capabilities : VL0 - VL7 Operational VLS : - Description : IB Subnet : infiniband-default Phy-profile : high-speed-ber Width reduction mode : Not supported Telemetry sampling : Disabled Telemetry threshold : Disabled Telemetry record : Disabled Telemetry threshold level : N/A bytes Rx: Bytes : 0 Packets : 0 Errors : 0 Symbol errors : 0 VL15 dropped packets : 0 Tx: Bytes : 0 Packets : 0 Wait : 0 Discarded packets : 0 </pre>
Related Commands	
Notes	<p>If a high power transceiver (e.g. LR4) is inserted to a port that does not support it, the link will not go up, and the following warning message is displayed: “Warning: High power transceiver is not supported” when the command “show interfaces ib” is run. For more information, please refer to see “High Power Transceivers”.</p>

10.4.5.15 show interfaces ib status

	show interfaces ib [<inf>] status Displays the status, speed and negotiation mode of the specified interface.	
Syntax Description	internal	Internal interfaces
	leaf-ports	filter to leaf-ports only
	inf	Interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.0500	
	3.4.1604	Updated example
	3.6.1002	Updated example

Example	<pre>switch (config) # show interfaces ib status Interface Description IB Subnet Speed Current line rate Logical port state Physical port state ----- IBl/1 infiniband-1 fdr 56.0 Gbps Active LinkUp IBl/2 infiniband-2 fdr 56.0 Gbps Active LinkUp IBl/3 infiniband-default - - Down Polling IBl/4 infiniband-default - - Down Polling IBl/5 infiniband-default - - Down Polling IBl/6 infiniband-default - - Down Polling IBl/7 infiniband-default - - Down Polling IBl/8 infiniband-default - - Down Polling IBl/9 infiniband-default - - Down Polling IBl/10 infiniband-default - - Down Polling IBl/11 infiniband-default - - Down Polling ...</pre>
Related Commands	
Notes	

10.4.5.16 show interfaces ib internal

	<pre>show interfaces ib internal [leaf spine] [<slot/module/port>] Displays running state for the internal ports of leafs or spines.</pre>	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.0500	
	3.9.3100	Updated field 'VL capabilities' to 'VL admin capabilities'
Example	<pre>switch (config) # show interfaces ib internal spine 1/1/4 IBl/1/4 state: Connected to slot/chip : 4/1 Connected to port : 19 Connected device active: 1 Error state : 0 Logical port state : Active Physical port state : LinkUp Current line rate : 56.0 Gbps Supported speeds : sdr, ddr, qdr, fdr10, fdr Speed : fdr Supported widths : 1X, 4X Width : 4X Max supported MTUs : 4096 MTU : 4096 VL admin capabilities : VL0 - VL7 Operational VLS : VL0 - VL7 Description : Phy-profile : high-speed-ber Width reduction mode : disabled</pre>	
Related Commands		

Notes	
-------	--

10.4.5.17 show interfaces ib internal capabilities

	show interfaces ib internal [leaf spine] [<slot/module/port>] capabilities Displays capabilities of internal leaf or spine interfaces.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.2.0500
Example	switch (config) # show interfaces ib internal leaf 1/1/26 capabilities IB1/1/26 LLR: FDR10, FDR,
Related Commands	
Notes	

10.4.5.18 show interfaces ib internal llr

	show interfaces ib internal [leaf spine] [<slot/module/port>] llr Displays LLR state of internal leaf or spine interfaces.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.2.0500 3.6.6000 Updated example
Example	switch (config) # show interfaces ib internal leaf 1/1/26 llr ----- Interface LLR status ----- IB1/1/26 Active
Related Commands	
Notes	

10.4.5.19 show interfaces ib internal status

	show interfaces ib internal [leaf spine] [<slot/module/port>] status Displays detailed running state of internal leaf or spine interfaces.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.2.0500

Example	<pre>switch (config) # show interfaces ib internal leaf 1/1/26 status</pre> <table border="1"> <thead> <tr> <th>Interface</th> <th>Description</th> <th>Speed</th> <th>Current line rate</th> <th>Logical port state</th> </tr> <tr> <th>Physical port state</th> <th colspan="4"></th> </tr> <tr> <th>-----</th> <th>-----</th> <th>-----</th> <th>-----</th> <th>-----</th> </tr> </thead> <tbody> <tr> <td>IB1/1/26</td> <td></td> <td>fdr</td> <td>56.0 Gbps</td> <td>Active</td> </tr> <tr> <td>LinkUp</td> <td colspan="4"></td> </tr> </tbody> </table>	Interface	Description	Speed	Current line rate	Logical port state	Physical port state					-----	-----	-----	-----	-----	IB1/1/26		fdr	56.0 Gbps	Active	LinkUp				
Interface	Description	Speed	Current line rate	Logical port state																						
Physical port state																										
-----	-----	-----	-----	-----																						
IB1/1/26		fdr	56.0 Gbps	Active																						
LinkUp																										
Related Commands																										
Notes																										

10.4.5.20 show interfaces ib transceiver

	show interfaces ib [<inf>] transceiver Displays the transceiver info.	
Syntax Description	inf	interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.10.6000	Updated example
Example	<pre>switch (config) # show interfaces ib 1/1/1 transceiver IB1/1/1 state: identifier : OSFP cable/module type : Optical module cable/module technology : 1310 nm EML infiniband speeds : FDR , EDR , HDR vendor : NVIDIA supported cable length : 100m SMF part number : MMS4X00-NL revision : A2 serial number : MT2216FI00088 FW version : 46.130.120</pre>	
Related Commands		
Notes	<ul style="list-style-type: none"> For a full list of the supported cables and transceivers, please refer to the LinkX™ Cables and Transceivers page If a high power transceiver (e.g. LR4) is used, it will be indicated in the field “cable/module type” 	

10.4.5.21 show interfaces ib transceiver diagnostics

	show interfaces ib [<inf>] transceiver diagnostics Displays cable channel monitoring and diagnostics info for this interface.	
Syntax Description	inf	Interface number: <slot>/<port>
Default	N/A	
Configuration Mode	Any command mode	
History	3.6.2002	
	3.6.6000	Updated example

<p>Example</p>	<pre> switch (config) # show interfaces ib transceiver diagnostics IB1/1 Transceiver Diagnostic Data: Message: No Diagnostic Data Available. Module is not DDMI capable IB1/3 Transceiver Diagnostic Data: Message: Non present module IB1/5 Transceiver Diagnostic Data: Temperature (-127C to +127C): Temperature : 28 C Hi Temp Alarm Thresh : 80 C Low Temp Alarm Thresh: -10 C Temperature Alarm : None Voltage (0 to 6.5535 V): Voltage : 3.28980 V Hi Volt Alarm Thresh : 3.50000 V Low Volt Alarm Thresh: 3.10000 V Voltage Alarm : None Tx Bias Current (0 to 131 mA): Ch1 Tx Current : 6.60000 mA Ch2 Tx Current : 6.60000 mA Ch3 Tx Current : 6.60000 mA Ch4 Tx Current : 6.60000 mA Hi Tx Crnt Alarm Thresh : 8.50000 mA Low Tx Crnt Alarm Thresh: 5.49200 mA Ch1 Tx Current Alarm : None Ch2 Tx Current Alarm : None Ch3 Tx Current Alarm : None Ch4 Tx Current Alarm : None Tx Power (0 mW to 6.5535 mW / 8.1647 dBm): Ch1 Tx Power : 1.01170 mW / 0.05052 dBm Ch2 Tx Power : 0.96240 mW / -0.16644 dBm Ch3 Tx Power : 0.95980 mW / -0.17819 dBm Ch4 Tx Power : 0.95800 mW / -0.18634 dBm Hi Tx Power Alarm Thresh : 3.46730 mW / 5.39991 dBm Low Tx Power Alarm Thresh: 0.07240 mW / -11.40261 dBm Ch1 Tx Power Alarm : None Ch2 Tx Power Alarm : None Ch3 Tx Power Alarm : None Ch4 Tx Power Alarm : None Rx Power (0 mW to 6.5535 mW / 8.1647 dBm): Ch1 Rx Power : 0.99160 mW / -0.03663 dBm Ch2 Rx Power : 1.08800 mW / 0.36629 dBm Ch3 Rx Power : 1.09810 mW / 0.40642 dBm Ch4 Rx Power : 0.97500 mW / -0.10995 dBm Hi Rx Power Alarm Thresh : 3.46730 mW / 5.39991 dBm Low Rx Power Alarm Thresh: 0.04670 mW / -13.30683 dBm Ch1 Rx Power Alarm : None Ch2 Rx Power Alarm : None Ch3 Rx Power Alarm : None Ch4 Rx Power Alarm : None Vendor Date Code (dd-mm-yyyy): 07-11-2016 </pre>
<p>Related Commands</p>	
<p>Notes</p>	<p>This example is for a QSFP transceiver</p>

10.4.5.22 show interfaces ib transceiver raw

	<p>show interfaces ib [<i><inf></i>] transceiver raw Displays cable info for this interface.</p>	
<p>Syntax Description</p>	<p>inf</p>	<p>interface number: <slot>/<port></p>
<p>Default</p>	<p>N/A</p>	
<p>Configuration Mode</p>	<p>Any command mode</p>	

History	3.6.1002
Example	<pre>switch (config) # show interfaces ib 1/7 transceiver raw IB1/7 raw transceiver data: I2C Address 0x50, Page 0, 0:255: 0000 0d 02 06 00 00 00 00 00 00 00 00 00 00 00 00 0010 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 0020 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 0030 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 0040 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 0050 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 0060 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 0070 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 0080 0d 00 23 08 00 00 00 00 00 00 00 00 05 8d 00 00 00 ..#. 0090 00 00 01 a0 4d 65 6c 6c 61 6e 6f 78 20 20 20 20Mellanox 00a0 20 20 20 20 0f 00 02 c9 4d 43 32 32 30 37 31 33MC220713 00b0 30 2d 30 30 41 20 20 20 41 33 02 03 05 00 46 66 0-00A A3....Ff 00c0 00 00 00 00 4d 54 31 32 32 37 56 53 30 30 36 34MT1227VS0064 00d0 32 20 20 20 31 32 30 37 30 38 20 20 00 00 00 e4 2 120708 00e0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00f0 00 00 00 00 00 00 00 00 00 00 00 02 00 00 30 00 00 I2C Address 0x50, Pages 1, 128:255: 0080 0d 02 06 00 00 00 00 00 00 00 00 00 00 00 00 0090 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00a0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00b0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00c0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00d0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00e0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00f0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00</pre>
Related Commands	
Notes	

10.5 Subnet Manager



The InfiniBand Subnet Manager (SM) is a centralized entity running in the switch. The SM discovers and configures all the InfiniBand fabric devices to enable traffic flow between those devices.

The SM applies network traffic related configurations such as Quality of Service (QoS), routing, and partitioning of the fabric devices. You can view and configure the Subnet Parameters (SM) via the CLI/WebUI menu. The embedded SM on the MLNX-OS can be used to manage fabrics up to 2048 nodes on x86 based systems.

The SM is used to discover and configure all the InfiniBand fabric devices to enable traffic flow between those devices.

- Subnet manager running via MLNX-OS does not support Dragonfly+ topology or combination of Fat-Tree and Dragonfly+ topologies.
- Subnet manager running via MLNX-OS does not support adaptive routing, fault routing (i.e., SHIELD or FRN), congestion control, SHIELD, and SHARP.

To enable Subnet Manager:

1. Enable Subnet Manager (disabled by default). Run:

```
switch (config) # ib smnode my-sm enable
```

2. (Optional) Set the priority for the Subnet Manager. Run:

```
switch (config) # ib smnode my-sm sm-priority <priority>
```

If rapid SM restarts are observed in what should be a quiet subnet, verify that all nodes running SM in the same management domain are in the same IB subnet. If they are not, fix the subnet.

10.5.1 Partitions

Partitioning enforces isolation among systems sharing an InfiniBand fabric. Partitioning is not related to boundaries established by subnets, switches, or routers. Rather, a partition describes a set of end nodes within the fabric that may communicate. Each port of an end node is a member of at least one partition and may be a member of multiple partitions. A partition manager (part of the SM) assigns partition keys (PKEYs) to each channel adapter port. Each PKEY represents a partition. Reception of an invalid PKEY causes the packet to be discarded. Switches and routers may optionally be used to enforce partitioning. In this case the partition manager programs the switch or router with PKEY information and when the switch or router detects a packet with an invalid PKEY, it discards the packet.

Fabric administration can assign certain Service Levels (SLs) for particular partitions. This allows the SM to isolate traffic flows between those partitions, and even if both partitions operate at the same QoS level, each partition can be guaranteed its fair share of bandwidth regardless of whether nodes in other partitions misbehave or are over subscribed.

The switch enables the configuration of partitions in an InfiniBand fabric.

The default partition is created by the SM unconditionally (whether it was defined or not).

10.5.1.1 Relationship with ib0 Interface

IP interface “ib0” is running under the default PKEY (0x7fff) and can be used for in-band management connectivity to the system.

10.5.1.2 Configuring Partition

The partitions configuration is applicable and to be used only when the SM is enabled and running on the system.

1. Create a partition. Run:

```
switch (config) # ib partition my-partition pkey 0x7ff2
```

2. Enter partition configuration mode. Run:

```
switch (config) # partition my-partition
```

```
switch (config partition name my-partition) #
```

3. Add partition members. Run:

```
switch (config partition my-partition) # member all
```

4. Verify the partition configuration. Run:

```
switch (config partition my-partition) # show ib partition
Default
  PKey      = 0x7FFF
  defmember = full
  ipoib     = yes
members
  GUID='ALL' member='full'
my-partition
  PKey      = 0x7ff2
members
  GUID='ALL' member='default'
```

10.5.2 Adaptive Routing

Adaptive routing (AR) allows optimizing data traffic flow. The InfiniBand protocol uses multiple paths between any two points. Thus, when unexpected traffic patterns cause some paths to be overloaded, AR can automatically move traffic to less congested paths according to the current temporal state of the network.

The embedded SM over the switch does not support configuring adaptive routing. To use this option in the fabric please use an external SM.

AR support is enabled by default on system profile “ib-single-switch”. To disable AR run either the command “system profile ib-no-adaptive-routing-single-switch” or “system profile ib” with no-adaptive-routing parameter.

The AR option needs to be enabled in the SM for it to take effect.

10.5.3 Scatter Ports

When assigning logical paths to physical links, the UpDn algorithm tries to map the same number of paths per link to maximize use of the available bandwidth. This balancing is done statically, without knowledge of actual workloads and traffic patterns. Path balancing decisions are made locally, at each switch, without assuming anything about the physical topology. The resulting path assignments may not be optimal for typical Clos/Fat Tree workloads.

A routing option called “scatter-ports” is available for MinHop and UpDn routing engines which instructs the routing algorithm to randomize the local assignments of paths to links, which often results in better link utilization. The scatter-ports option requires an integer argument, which is the seed for the random number generator. It is recommended to use a prime number for the seed; a seed of zero turns off randomization.

10.5.4 GUID Routing Order

GUID routing order list allows managing the order in which the SM processes the destination LIDs in the calculations of output port as part of MinHop or Up/Down routing algorithms only.

The order of GUID appearance is important as destinations corresponding to GUIDs appearing earlier in the routing list get precedence during the routing calculations over other destinations in the fabric. This can improve load balancing towards a specific set of end ports (e.g. storage nodes or other service nodes requiring high throughput).

If scatter-ports (randomization of the output port) option is set to non-zero, `guid-routing-order-no-scatter` defines whether or not a randomization should be applied to the destinations GUIDs mentioned in GUID routing order list.

10.5.5 Bulk Update Mode

Bulk update mode allows users to set multiple IB SM configurations without applying them until bulk mode is disabled.

When bulk update is disabled (default situation) every SM configuration is applied immediately. When bulk is enabled, all SM configuration is saved internally and is not applied until this mode is disabled.

Bulk mode is a non-persistent state. That is, if the switch is restarted, it boots up with this mode disabled, and all the configuration changes which are saved before system restart are applied.

Show commands convey every configuration change even if it is not applied yet.

10.5.6 SM Commands



- [10.5.6.1 General](#)
 - [10.5.6.1.1 `ib sm`](#)
 - [10.5.6.1.2 `ib sm accum-log-file`](#)
 - [10.5.6.1.3 `ib sm allow-both-pkeys`](#)
 - [10.5.6.1.4 `ib sm babbling-policy`](#)
 - [10.5.6.1.5 `ib sm connect-roots`](#)
 - [10.5.6.1.6 `ib sm calculate-missing-routes`](#)
 - [10.5.6.1.7 `ib sm drop-event-subscriptions`](#)
 - [10.5.6.1.8 `ib sm enable-quirks`](#)
 - [10.5.6.1.9 `ib sm exit-on-fatal`](#)
 - [10.5.6.1.10 `ib sm force-link-speed`](#)
 - [10.5.6.1.11 `ib sm force-log-flush`](#)
 - [10.5.6.1.12 `ib sm guid2lid-cache`](#)
 - [10.5.6.1.13 `ib sm honor-partitions`](#)
 - [10.5.6.1.14 `ib sm hoq-lifetime`](#)
 - [10.5.6.1.15 `ib sm ignore-other-sm`](#)
 - [10.5.6.1.16 `ib sm ipv6-nsm`](#)
 - [10.5.6.1.17 `ib sm lash`](#)
 - [10.5.6.1.18 `ib sm leafhoq-lifetime`](#)
 - [10.5.6.1.19 `ib sm leafvl-stalls`](#)
 - [10.5.6.1.20 `ib sm lmc`](#)
 - [10.5.6.1.21 `ib sm lmc-esp0`](#)

- [10.5.6.1.22 ib sm log-flags](#)
- [10.5.6.1.23 ib sm log-max-size](#)
- [10.5.6.1.24 ib sm max-op-vls](#)
- [10.5.6.1.25 ib sm max-reply-time](#)
- [10.5.6.1.26 ib sm max-reverse-hops](#)
- [10.5.6.1.27 ib sm aguid_default_hop_limit](#)
- [10.5.6.1.28 ib sm max-wire-smpps2](#)
- [10.5.6.1.29 ib sm m-key](#)
- [10.5.6.1.30 ib sm mkey-lease](#)
- [10.5.6.1.31 ib sm mkey-lookup](#)
- [10.5.6.1.32 ib sm mkey-protect-level](#)
- [10.5.6.1.33 ib sm msgfifo-timeout](#)
- [10.5.6.1.34 ib sm multicast](#)
- [10.5.6.1.35 ib sm no-client-rereg](#)
- [10.5.6.1.36 ib sm overrun-trigger](#)
- [10.5.6.1.37 ib sm packet-life-time](#)
- [10.5.6.1.38 ib sm phy-err-trigger](#)
- [10.5.6.1.39 ib sm polling-retries](#)
- [10.5.6.1.40 ib sm port-prof-switch](#)
- [10.5.6.1.41 ib sm reassign-lids](#)
- [10.5.6.1.42 ib sm reset-config](#)
- [10.5.6.1.43 ib sm root-guid](#)
- [10.5.6.1.44 ib sm routing-engines](#)
- [10.5.6.1.45 ib sm rtr-aguid-enable](#)
- [10.5.6.1.46 ib sm rtr-pr-flow-label](#)
- [10.5.6.1.47 ib sm rtr-pr-mtu](#)
- [10.5.6.1.48 ib sm rtr-pr-rate](#)
- [10.5.6.1.49 ib sm rtr-pr-sl](#)
- [10.5.6.1.50 ib sm rtr-pr-tclass](#)
- [10.5.6.1.51 ib sm sa-key](#)
- [10.5.6.1.52 ib sm single-thread](#)
- [10.5.6.1.53 ib sm sm-inactive](#)
- [10.5.6.1.54 ib sm sm-key](#)
- [10.5.6.1.55 ib sm sm-priority](#)
- [10.5.6.1.56 ib sm sm-sl](#)
- [10.5.6.1.57 ib sm sminfo-poll-time](#)
- [10.5.6.1.58 ib sm subnet-prefix](#)
- [10.5.6.1.59 ib sm subnet-prefix-override](#)
- [10.5.6.1.60 ib sm max-smpps-timeout](#)
- [10.5.6.1.61 ib sm subnet-timeout](#)
- [10.5.6.1.62 ib sm sweep-interval](#)
- [10.5.6.1.63 ib sm sweep-on-trap](#)
- [10.5.6.1.64 ib sm transaction-retries](#)
- [10.5.6.1.65 ib sm use-heavy-sweeps](#)
- [10.5.6.1.66 ib sm use-ucast-cache](#)
- [10.5.6.1.67 ib sm vl-stalls](#)
- [10.5.6.1.68 ib sm virt](#)

- [10.5.6.1.69 ib sm virt-default-hop-limit](#)
- [10.5.6.1.70 ib sm virt-max-ports-in-process](#)
- [10.5.6.2 Show](#)
 - [10.5.6.2.1 show ib sm](#)
 - [10.5.6.2.2 show ib sm accum-log-file](#)
 - [10.5.6.2.3 show ib sm babbling-policy](#)
 - [10.5.6.2.4 show ib sm calculate-missing-routes](#)
 - [10.5.6.2.5 show ib sm connect-roots](#)
 - [10.5.6.2.6 show ib sm enable-quirks](#)
 - [10.5.6.2.7 show ib sm exit-on-fatal](#)
 - [10.5.6.2.8 show ib sm fdr10](#)
 - [10.5.6.2.9 show ib sm force-link-speed](#)
 - [10.5.6.2.10 show ib sm force-link-speed-ext](#)
 - [10.5.6.2.11 show ib sm force-log-flush](#)
 - [10.5.6.2.12 show ib sm guid2lid-cache](#)
 - [10.5.6.2.13 show ib sm honor-partitions](#)
 - [10.5.6.2.14 show ib sm hoq-lifetime](#)
 - [10.5.6.2.15 show ib sm ignore-other-sm](#)
 - [10.5.6.2.16 show ib sm ipv6-nsm](#)
 - [10.5.6.2.17 show ib sm lash](#)
 - [10.5.6.2.18 show ib sm leafhoq-lifetime](#)
 - [10.5.6.2.19 show ib sm leafvl-stalls](#)
 - [10.5.6.2.20 show ib sm lmc](#)
 - [10.5.6.2.21 show ib sm lmc-esp0](#)
 - [10.5.6.2.22 show ib sm log](#)
 - [10.5.6.2.23 show ib sm log-flags](#)
 - [10.5.6.2.24 show ib sm log-max-size](#)
 - [10.5.6.2.25 show ib sm max-op-vls](#)
 - [10.5.6.2.26 show ib sm max-ports](#)
 - [10.5.6.2.27 show ib sm max-reply-time](#)
 - [10.5.6.2.28 show ib sm max-reverse-hops](#)
 - [10.5.6.2.29 show ib sm aguid-default-hop-limit](#)
 - [10.5.6.2.30 show ib sm max-wire-smpps](#)
 - [10.5.6.2.31 show ib sm max-wire-smpps2](#)
 - [10.5.6.2.32 show ib sm mkey-lease](#)
 - [10.5.6.2.33 show ib sm m-key](#)
 - [10.5.6.2.34 show ib sm mkey-lease](#)
 - [10.5.6.2.35 show ib sm mkey-lookup](#)
 - [10.5.6.2.36 show ib sm mkey-protect-level](#)
 - [10.5.6.2.37 show ib sm msgfifo-timeout](#)
 - [10.5.6.2.38 show ib sm multicast](#)
 - [10.5.6.2.39 show ib sm no-client-rereg](#)
 - [10.5.6.2.40 show ib sm overrun-trigger](#)
 - [10.5.6.2.41 show ib sm packet-life-time](#)
 - [10.5.6.2.42 show ib sm phy-err-trigger](#)
 - [10.5.6.2.43 show ib sm polling-retries](#)
 - [10.5.6.2.44 show ib sm port-prof-switch](#)

- [10.5.6.2.45 show ib sm reassign-lids](#)
- [10.5.6.2.46 show ib sm root-guid](#)
- [10.5.6.2.47 show ib sm routing-engines](#)
- [10.5.6.2.48 show ib sm routing-info](#)
- [10.5.6.2.49 show ib sm rtr-aguid-enable](#)
- [10.5.6.2.50 show ib sm rtr-pr-flow-label](#)
- [10.5.6.2.51 show ib sm rtr-pr-mtu](#)
- [10.5.6.2.52 show ib sm rtr-pr-rate](#)
- [10.5.6.2.53 show ib sm rtr-pr-sl](#)
- [10.5.6.2.54 show ib sm sa-key](#)
- [10.5.6.2.55 show ib sm single-thread](#)
- [10.5.6.2.56 show ib sm sm-inactive](#)
- [10.5.6.2.57 show ib sm sm-key](#)
- [10.5.6.2.58 show ib sm sm-priority](#)
- [10.5.6.2.59 show ib sm sm-sl](#)
- [10.5.6.2.60 show ib sm sminfo-poll-time](#)
- [10.5.6.2.61 show ib sm subnet-prefix](#)
- [10.5.6.2.62 show ib sm subnet-prefix-override](#)
- [10.5.6.2.63 show ib sm subnet-timeout](#)
- [10.5.6.2.64 show ib sm sweep-interval](#)
- [10.5.6.2.65 show ib sm sweep-on-trap](#)
- [10.5.6.2.66 show ib sm transaction-retries](#)
- [10.5.6.2.67 show ib sm use-heavy-sweeps](#)
- [10.5.6.2.68 show ib sm use-ucast-cache](#)
- [10.5.6.2.69 show ib sm version](#)
- [10.5.6.2.70 show ib sm virt-default-hop-limit](#)
- [10.5.6.2.71 show ib sm virt-max-ports-in-process](#)
- [10.5.6.2.72 show ib sm vl-stalls](#)
- [10.5.6.3 Partitions](#)
 - [10.5.6.3.1 ib partition](#)
 - [10.5.6.3.2 pkey](#)
 - [10.5.6.3.3 defmember](#)
 - [10.5.6.3.4 member](#)
 - [10.5.6.3.5 ipoib](#)
 - [10.5.6.3.6 mtu](#)
 - [10.5.6.3.7 rate](#)
 - [10.5.6.3.8 scope](#)
 - [10.5.6.3.9 sl](#)
 - [10.5.6.3.10 show ib partition](#)
- [10.5.6.4 Quality of Service \(SM\)](#)
 - [10.5.6.4.1 ib baseqos <port-type> high-limit](#)
 - [10.5.6.4.2 ib baseqos max-vls](#)
 - [10.5.6.4.3 ib baseqos sl2vl](#)
 - [10.5.6.4.4 ib baseqos vlarb-high](#)
 - [10.5.6.4.5 ib baseqos <port-type> vlarb-low <value>](#)
 - [10.5.6.4.6 ib baseqos reset-config](#)
 - [10.5.6.4.7 show ib baseqos](#)

- [10.5.6.4.8 ib qos](#)
- [10.5.6.4.9 ib qos level](#)
- [10.5.6.4.10 ib qos match-rule](#)
- [10.5.6.4.11 ib qos port-group](#)
- [10.5.6.4.12 ib qos ulp any](#)
- [10.5.6.4.13 ib qos ulp ipoib](#)
- [10.5.6.4.14 ib qos ulp <protocol-type>](#)
- [10.5.6.4.15 ib qos ulp srp](#)
- [10.5.6.4.16 show ib qos](#)
- [10.5.6.5 Scatter Ports](#)
 - [10.5.6.5.1 ib sm scatter-ports](#)
 - [10.5.6.5.2 show ib sm scatter-ports](#)
- [10.5.6.6 GUID Routing Order](#)
 - [10.5.6.6.1 ib sm guid-routing-order add](#)
 - [10.5.6.6.2 ib sm guid-routing-order delete](#)
 - [10.5.6.6.3 ib sm guid-routing-order move](#)
 - [10.5.6.6.4 ib sm guid-routing-order move-down](#)
 - [10.5.6.6.5 ib sm guid-routing-order move-up](#)
 - [10.5.6.6.6 no ib sm guid-routing-order](#)
 - [10.5.6.6.7 ib sm guid-routing-order-no-scatter](#)
 - [10.5.6.6.8 show ib sm guid-routing-order](#)
 - [10.5.6.6.9 show ib sm guid-routing-order-no-scatter](#)
- [10.5.6.7 Bulk Update Mode](#)
 - [10.5.6.7.1 ib sm bulk-update enable](#)
 - [10.5.6.7.2 show ib sm bulk-update](#)
- [10.5.6.8 ibdiagnet](#)
 - [10.5.6.8.0 ibdiagnet](#)
 - [10.5.6.8.2 show ibdiagnet](#)
 - [10.5.6.8.3 file ibdiagnet upload](#)
 - [10.5.6.8.4 file ibdiagnet delete](#)

10.5.6.1 General

10.5.6.1.1 ib sm

	ib sm no ib sm Enables the SM on this node. The no form of the command disables the SM on this node.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.1.0000
Example	switch (config) # ib sm
Related Commands	show ib sm

Notes	
-------	--

10.5.6.1.2 `ib sm accum-log-file`

	<code>ib sm accum-log-file</code> <code>no ib sm accum-log-file</code> Adds SM log entries at the end of the current log. The no form of the command overwrites SM log file on every restart.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config
History	3.1.0000
Example	<code>switch (config) # ib sm accum-log-file</code>
Related Commands	<code>show ib sm accum-log-file</code>
Notes	

10.5.6.1.3 `ib sm allow-both-pkeys`

	<code>ib sm allow-both-pkeys</code> <code>no ib sm allow-both-pkeys</code> Enables having both full and limited membership on the same partition. The no form of the command disables having both full and limited membership on the same partition.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.4.1100
Example	<code>switch (config) # ib sm allow-both-pkeys</code>
Related Commands	<code>defmember</code> <code>member</code>
Notes	

10.5.6.1.4 `ib sm babbling-policy`

	<code>ib sm babbling-policy</code> <code>no ib sm babbling-policy</code> Enables the SM to disable babbling ports (i.e., generating frequent traps). The no form of the command disables the SM babbling policy.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.1.0000

Example	<code>switch (config) # no ib sm babbling-policy</code>
Related Commands	<code>show ib sm babbling-policy</code>
Notes	If the babbling policy is enabled, and decides to close a babbling interface (one which sends 129,130,131 traps, for example), the SM disables the port.

10.5.6.1.5 ib sm connect-roots

	<code>ib sm connect-roots</code> <code>no ib sm connect-roots</code> Forces the routing engine to make connectivity between root switches. The no form of the command disables logical LID path between root switches.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config
History	3.1.0000
Example	<code>switch (config) # ib sm connect-roots</code>
Related Commands	<code>show ib sm connect-roots</code>
Notes	<ul style="list-style-type: none"> • This command is relevant only for ‘updn’ and ‘ftree’ algorithm (refer to ‘ib sm routing-engines’ command) • This option enforces routing engines (up/down and fat-tree) to make connectivity between root switches and in this way to be fully IBA compliant. This may violate the “deadlock-free” status of the algorithm. Hence, it is recommended to use the command carefully.

10.5.6.1.6 ib sm calculate-missing-routes

	<code>ib sm calculate-missing-routes</code> <code>no ib sm calculate-missing-routes</code> Enables SM to find and recalculate missing routes without creating credit-loops The no form of the command disables SM to find and recalculate missing routes without creating credit-loops
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.8.2000
Example	<pre>switch (config) # ib sm calculate-missing-routes switch (config) # show ib sm calculate-missing-routes ib sm calculate-missing-routes: enabled switch (config) # no ib sm calculate-missing-routes switch (config) # show ib sm calculate-missing-routes ib sm calculate-missing-routes: disabled</pre>
Related Commands	<code>Show ib sm calculate-missing-routes</code>
Notes	

10.5.6.1.7 ib sm drop-event-subscriptions

	<code>ib sm drop-event-subscriptions</code> <code>no ib sm drop-event-subscriptions</code> Configures IB SM to drop interface subscribe or unsubscribe events. The no form of the command resets this parameter to its default value.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.4.2008
Example	<code>switch (config) # ib sm drop-event-subscriptions</code>
Related Commands	
Notes	

10.5.6.1.8 ib sm enable-quirks

Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.1.0000
Example	<code>switch (config) # ib sm enable-quirks</code>
Related Commands	<code>show ib sm enable-quirks</code>
Notes	

10.5.6.1.9 ib sm exit-on-fatal

	<code>ib sm exit-on-fatal</code> <code>no ib sm exit-on-fatal</code> Enables the SM to exit upon fatal initialization errors. The no form of the command disables the SM from exiting upon fatal initialization errors.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config
History	3.1.0000
Example	<code>switch (config) # ib sm exit-on-fatal</code>
Related Commands	<code>show ib sm exit-on-fatal</code>
Notes	

10.5.6.1.10 ib sm force-link-speed

	ib sm force-link-speed <speed-options> no ib sm force-link-speed Defines the SM behavior for PortInfo:LinkSpeedEnabled, PortInfo:LinkSpeedExtEnabled and MLNX ExtendedPortInfo on the switch ports. The no form of the command resets this parameter to its default.	
Syntax Description	speed-options	The following options are available: <ul style="list-style-type: none"> • sdr—10.0 Gb/s rate on 4 lane width • ddr—20.0 Gb/s rate on 4 lane width • qdr—40.0 Gb/s rate on 4 lane width • fdr10—40.0 Gb/s rate on 4 lane width • fdr—56.0 Gb/s rate on 4 lane width • edr—100.0 Gb/s rate on 4 lane width
Default	Set to PortInfo:LinkSpeedExtSupported	
Configuration Mode	config	
History	3.1.0000	
	3.4.1604	Updated Syntax Description, Example, and Notes
Example	switch (config) # ib sm force-link-speed sdr ddr qdr fdr10	
Related Commands	show ib sm force-link-speed show ib sm force-link-speed-ext show ib sm fdr10	
Notes	<ul style="list-style-type: none"> • The following options, as defined in InfiniBand Specification 1.2.1 section 14.2.5.6, table 145 “PortInfo” • This command updates force-link-speed, force-link-speed ext and fdr10 which are open sm parameters • This command is backwards compatible so old configuration file containing this command with the old form (with legal bit mask) are still supported • If the speed-options list does not include SDR speed, it is configured automatically • Configuring more than one speed is possible by typing in consecutive speed names separated by spaces 	

10.5.6.1.11 ib sm force-log-flush

	ib sm force-log-flush no ib sm force-log-flush Forces every log message generated to be flushed. The no form of the command does not force a flush after every log write.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm force-log-flush	
Related Commands	show ib sm force-log-flush	
Notes		

10.5.6.1.12 `ib sm guid2lid-cache`

	<code>ib sm guid2lid-cache</code> <code>no ib sm guid2lid-cache</code> Allows SM to use cached GUID-to-lid mapping data. When enabled, the SM honors the cached GUID-to-lid mapping information if: <ul style="list-style-type: none"> • It exists • It is valid • <code>sm_reassign_lids</code> is disabled The no form of the command disallows use of cached GUID-to-lid mapping data.
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.1.0000
Example	<code>switch (config) # ib sm guid2lid-cache</code>
Related Commands	<code>show ib sm guid2lid-cache</code>
Notes	

10.5.6.1.13 `ib sm honor-partitions`

	<code>ib sm honor-partitions</code> <code>no ib sm honor-partitions</code> Sets the <code>no_partition_enforcement</code> flag to 0. This setting controls global support for partitioning in the subnet. The no form of the command disables subnet partition support.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config
History	3.1.0000
Example	<code>switch (config) # no ib sm honor-partitions</code>
Related Commands	<code>show ib sm honor-partitions</code>
Notes	<ul style="list-style-type: none"> • If partitioning is disabled (<code>no_partition_enforcement=1</code>), then no named partitions can be enabled • If partitioning is enabled globally, the <code>no_partition_enforcement</code> changes from 1 to 0, and all defined partitions with state enabled are instantiated • If partitioning is globally disabled, all partitions are removed from the subnet, but the state (enabled or disabled) associated with defined partitions is not modified

10.5.6.1.14 `ib sm hoq-lifetime`

	<code>ib sm hoq-lifetime <time></code> Sets the maximum time a frame can wait at the head of a switch-to-switch port queue before it is dropped.
--	---

Syntax Description	time	The time is 4.096 uS * 2time. The range of time is 0 to 20. A time of 20 means infinite, and the default value is 18 which translates to about 1 second.
Default	0x12 (~ 1 second)	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm hoq-lifetime 15	
Related Commands	show ib sm hoq-lifetime	
Notes		

10.5.6.1.15 ib sm ignore-other-sm

	ib sm ignore-other-sm no ib sm ignore-other-sm Ignores all the rules governing SM elections and attempts to manage the fabric. The no form of the command does not allow the SM to manage fabric if it loses the election.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm ignore-other-sm	
Related Commands	show ib sm ignore-other-sm	
Notes		

10.5.6.1.16 ib sm ipv6-nsm

	ib sm ipv6-nsm no ib sm ipv6-nsm Consolidates IPv6 SNM group joins to 1 MC group per-MGID PKEY. The no form of the command disables the consolidation of IPv6 SNM.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm ipv6-nsm	
Related Commands	show ib sm ipv6-nsm	
Notes		

10.5.6.1.17 ib sm lash

	<pre>ib sm lash {do-mesh-analysis start-vl <vl-value>} no ib sm lash do-mesh-analysis</pre> <p>Modifies “lash” routing method parameters. The no form of the command disables SM “lash” routing for mesh analysis.</p>	
Syntax Description	do-mesh-analysis	Enables SM “lash” routing for mesh analysis
	start-vl <vl-value>	Configures the starting VL for SM “lash” routing for mesh analysis (assuming that lash routing is enabled)
Default	do-mesh-analysis: Disabled start-vl: 0	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm lash do-mesh-analysis	
Related Commands	show ib sm lash do-mesh-analysis	
Notes		

10.5.6.1.18 ib sm leafhoq-lifetime

	<pre>ib sm leafhoq-lifetime <time></pre> <p>Sets the maximum time a frame can wait at the head of a switch-to-CA_or_Router port queue before it is dropped.</p>	
Syntax Description	time	The time is $4.096 \mu\text{s} * 2\text{time}$. The range of time is 0 to 20. A time of 20 means infinite, and the default value is 16 which translates to about 268 millisecond.
Default	0x10 (about 268 mS)	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm leafhoq-lifetime 8	
Related Commands	show ib sm leafhoq-lifetime	
Notes		

10.5.6.1.19 ib sm leafvl-stalls

	<pre>ib sm leafvl-stalls <count></pre> <p>Sets the number of sequential frame drops that cause a switch-to-CA_or_Router port to enter the VLStalled state.</p>	
Syntax Description	count	Range: 1-255
Default	7	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm leafvl-stalls 3	

Related Commands	show ib sm leafvl-stalls
Notes	

10.5.6.1.20 ib sm lmc

	ib sm lmc <mask> Sets the LID Mask Control (LMC) value to be used on this subnet.	
Syntax Description	mask	Range: 0-7
Default	The default value is 0, which means that every port has exactly one unique LID.	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm lmc 7	
Related Commands	show ib sm lmc	
Notes		

10.5.6.1.21 ib sm lmc-esp0

	ib sm lmc-esp0 no ib sm lmc-esp0 Sets the LMC for the subnet to be used for Enhanced Switch Port 0. The no form of the command resets this parameter to its default.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm lmc-esp0	
Related Commands	show ib sm lmc-esp0	
Notes		

10.5.6.1.22 ib sm log-flags

	ib sm log-flags [all] [debug] [error] [frames] [funcs] [info] [none] [routing] [verbose] no ib sm log-flags Controls what messages the SM logs. The no form of the command indicates to the SM not to run on this node.	
Syntax Description	all	Turns on all the flags that follow (error info verbose debug funcs frames routing).
	debug	Logs diagnostic messages, high volume.
	error	Logs error messages.
	frames	Logs all SMP and GMP frames.
	funcs	Logs function entry/exit, very high volume.

	info	Logs basic messages, low volume.
	none	Turns off all logging flags.
	routing	Logs FDB routing information.
	verbose	Logs interesting stuff, moderate volume.
Default	0x3 (error, info)	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm log-flags error verbose funcs frames	
Related Commands	show ib sm log-flags	
Notes	<ul style="list-style-type: none"> • Every execution of this command replaces the current logging flags • The options “all” and “none” must be specified as the only parameter 	

10.5.6.1.23 ib sm log-max-size

	ib sm log-max-size <size> Sets the maximum size of the log file to be <size> megabytes.	
Syntax Description	size	Range: 1-60
Default	20 MBytes	
Configuration Mode	config	
History	3.1.0000	
	3.5.1000	Updated Syntax Description, and Default
Example	switch (config) # ib sm log-max-size 50	
Related Commands	show ib sm log-max-size	
Notes	<ul style="list-style-type: none"> • The log file “opensm_<switch_name>.log” is rotated when it exceeds the configured maximum file size up to 5 compressed files • When the log gets to the maximum size, or system storage fills up, the current log is deleted and messages start accumulating • To successfully upgrade from a version prior to 3.5.1000, this parameter must be set to a value in the range specified in the syntax descriptio 	

10.5.6.1.24 ib sm max-op-vls

	ib sm max-op-vls <count> Sets the maximum number of VLS supported on this subnet.	
Syntax Description	count	Possible values: 1, 2, 4, 8, or 15
Default	4	
Configuration Mode	config	
History	3.1.0000	
	3.10.1000	Updated default value from 15 to 4
Example	switch (config) # ib sm max-op-vls 4	
Related Commands	show ib sm max-op-vls	

Notes	
-------	--

10.5.6.1.25 ib sm max-reply-time

	<code>ib sm max-reply-time <time></code> Sets the maximum time the SM waits for a reply before the transaction times out.	
Syntax Description	time	Must be an integer (in milliseconds)
Default	200 milliseconds	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # ib sm max-reply-time 500</pre>	
Related Commands	show sm max-reply-time	
Notes		

10.5.6.1.26 ib sm max-reverse-hops

	<code>ib sm max-reverse-hops <max-reverse-hops></code> Sets the maximum number of hops from the top switch to an I/O node.	
Syntax Description	N/A	
Default	0 hops	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # ib sm max-reverse-hops 500</pre>	
Related Commands	show ib sm max-reverse-hops	
Notes		

10.5.6.1.27 ib sm aguid_default_hop_limit

	<code>ib sm aguid-default-hop-limit <count></code> <code>no ib sm aguid-default-hop-limit</code> Configures the default value for hop limit returned in path records where either the source or destination are alias an GUID. The no form of the command resets the count to its default value.	
Syntax Description	count	Number of concurrent management packets (must be an integer)
Default	1	
Configuration Mode	config	
History	3.6.6102	
Example	<pre>switch (config) # ib sm aguid-default-hop-limit 3</pre>	
Related Commands	show ib sm aguid-default-hop-limit	
Notes		

10.5.6.1.28 ib sm max-wire-smpls2

	ib sm max-wire-smpls2 <count> no Sets the maximal timeout based outstanding SM management packets. The no form of the command resets the max-wire-smpls2 to its initial value.	
Syntax Description	count	Number of concurrent management packets. The value must be an integer.
Default	4	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm max-wire-smpls 8	
Related Commands	show ib sm max-wire-smpls2	
Notes		

10.5.6.1.29 ib sm m-key

	ib sm m-key <mkey> no ib sm m-key Configures the MKey used by the SM. The no form of the command resets the MKey configuration to its default value.	
Syntax Description	mkey	64-bit MKey
Default	00:00:00:00:00:00:00:00	
Configuration Mode	config	
History	3.1.0000	
	3.6.2002	Added no form of the command
	3.7.0000	Added note
Example	switch (config) # ib sm m-key 11:33:55:77:99:aa:cc:ee	
Related Commands	ib sm mkey-lease ib sm mkey-lookup ib sm mkey-protect-level show ib sm m-key show ib sm mkey-lease	
Notes	<ul style="list-style-type: none"> • All nodes in the subnet may have to be reset or power-cycled after altering the SM MKey configuration • Fabric inspector, and many standalone InfiniBand utilities, may not function on subnets with a non-default MKey. 	

10.5.6.1.30 ib sm mkey-lease

	ib sm mkey-lease <time> no ib sm mkey-lease Configures the lease period used when MKey is non-zero. The no form of the command resets this value to its default.	
--	---	--

Syntax Description	time	MKey lease period in seconds Range: 0-65535; 0=unlimited
Default	0	
Configuration Mode	config	
History	3.6.2002	
Example	switch (config) # ib sm mkey-lease 660	
Related Commands	show ib sm mkey-lease	
Notes		

10.5.6.1.31 ib sm mkey-lookup

	ib sm mkey-lookup no ib sm mkey-lookup Enables using a file cache (guid2mkey) to resolve unknown node MKey. The no form of the command disables using a file cache to resolve unknown node MKey and the configured MKey is used for all ports.	
Syntax Description	N/A	
Default	Enabled	
Configuration Mode	config	
History	3.6.2002	
Example	switch (config) # ib sm mkey-lookup	
Related Commands	show ib sm mkey-lookup	
Notes	MKey lookup is a boolean value that controls how the SM finds the MKey of ports	

10.5.6.1.32 ib sm mkey-protect-level

	ib sm mkey-protect-level <level> no ib sm mkey-protect-level Controls what data is returned to a get_PortInfo MAD request when the MKey in the request does not match the MKey on the port. The no form of the command resets the parameter to its default value.	
Syntax Description	level	<ul style="list-style-type: none"> • 0—when PortInfo is “read”, the actual MKey is returned in port info data • 1—when PortInfo is “read”, and the MKey in the MAD does not match the MKey on the port, the MKey value in the returned PortInfo data is set to 0 • 2—when PortInfo is “read”, and the MKey in the MAD does not match the MKey on the port, no data is returned
Default	0	
Configuration Mode	config	
History	3.6.2002	
Example	switch (config) # ib sm mkey-protect-level 0	
Related Commands	show ib sm mkey-protect-level	
Notes		

10.5.6.1.33 `ib sm msgfifo-timeout`

	<code>ib sm msgfifo-timeout <time></code> Sets the time value to be used by the subnet administrator to control when a BUSY status is returned to a client.	
Syntax Description	<code>time</code>	In milliseconds
Default	10 seconds	
Configuration Mode	config	
History	3.1.0000	
Example	<code>switch (config) # ib sm msgfifo-timeout 50000</code>	
Related Commands	show ib sm msgfifo-timeout	
Notes	If there is more than one message in the SA queue, and it has been there longer than time milliseconds, all additional incoming requests are immediately replied to with BUSY status.	

10.5.6.1.34 `ib sm multicast`

	<code>ib sm multicast</code> <code>no ib sm multicast</code> Enables the SM to support multicasts on the fabric. The no form of the command disables the SM from supporting multicasts on the fabric.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config	
History	3.1.0000	
Example	<code>switch (config) # ib sm multicast</code>	
Related Commands	show ib sm multicast	
Notes		

10.5.6.1.35 `ib sm no-client-rereg`

	<code>ib sm no-client-rereg</code> <code>no ib sm no-client-rereg</code> Enables client re-registration requests. The no form of the command disables client re-registration requests.	
Syntax Description	N/A	
Default	disable	
Configuration Mode	config	
History	3.1.0000	
Example	<code>switch (config) # ib sm no-client-rereg</code>	
Related Commands	show ib sm no-client-rereg	
Notes		

10.5.6.1.36 `ib sm overrun-trigger`

	<code>ib sm overrun-trigger <count></code> Enables SMA to generate standard InfiniBand trap number 130 when the number of local buffer overrun errors equals the count value, and the port's SMA supports traps.	
Syntax Description	count	Range: 0-255
Default	8	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm overrun-trigger 3	
Related Commands	show ib sm overrun-trigger	
Notes	Refer to the InfiniBand Architecture Specification V1 r1.2.1, section 14.2.5.1 table 131: Traps.	

10.5.6.1.37 `ib sm packet-life-time`

	<code>ib sm packet-life-time <time></code> Sets the maximum time a frame can live in a switch.	
Syntax Description	time	The time is $4.096 \mu\text{S} * 2^{*}\langle\text{time}\rangle$. Range: 0-20. A time of 20 means infinite. The value 0x14 disables this mechanism.
Default	0x12 (about 1 second)	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm packet-life-time 20	
Related Commands	show ib sm packet-life-time	
Notes		

10.5.6.1.38 `ib sm phy-err-trigger`

	<code>ib sm phy-err-trigger <count></code> Enables SMA to generate trap 129 when the number of local link integrity errors equals the <count> value, and the port's SMA supports traps.	
Syntax Description	count	Range: 0-255
Default	8	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm phy-err-trigger 5	
Related Commands	show ib sm phy-err-trigger	
Notes		

10.5.6.1.39 `ib sm polling-retries`

	<code>ib sm polling-retries <value></code> This variable defines the number of consecutive times an active SM must fail to respond before it is declared dead.	
Syntax Description	value	Must be an integer
Default	4	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # ib sm polling-retries 8</pre>	
Related Commands	show ib sm polling-retries	
Notes	The time between when the active SM fails and the time this SM declares it dead is: (sm_sminfo_polling_timeout * value) milliseconds.	

10.5.6.1.40 `ib sm port-prof-switch`

	<code>ib sm port-prof-switch</code> <code>no ib sm port-prof-switch</code> Enables the counting of adapters, routers, and switches routed through links. The no form of the command disables the counting of adapters, routers, and switches routed through links.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # ib sm port-prof-switch</pre>	
Related Commands	show ib sm port-prof-switch	
Notes		

10.5.6.1.41 `ib sm reassign-lids`

	<code>ib sm reassign-lids</code> <code>no ib sm reassign-lids</code> Controls the ability of the SM to reassign LIDs to nodes it finds already configured with a valid LID. The no form of the command disables the SM from reassigning LIDs to nodes it finds already configured with a valid LID.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # ib sm reassign-lids</pre>	
Related Commands	show ib sm reassign-lids	

Notes	<ul style="list-style-type: none"> • If enabled (ib sm reassign-lids), the SM can, but is not required to, reassign the LID on a node with a pre-configured LID • If disabled (no ib sm reassign-lids), the SM does not reassign LIDs • There are times when the SM is required to reassign LIDs or the fabric cannot be brought to a stable state, or a fabric option (like LMC) can not be fully applied
-------	---

10.5.6.1.42 ib sm reset-config

	ib sm reset-config Resets all SM configuration options to defaults.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm reset-config	
Related Commands		
Notes		

10.5.6.1.43 ib sm root-guid

	ib sm root-guid <guid> no ib sm root-guid <guid> Adds a root GUID for the SM. The no form of the command removes the GUID from the root GUID list.	
Syntax Description	guid	The root GUID number in hexadecimal notation
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config)# ib sm root-guid aa:bb:00:11:22:33:44:55	
Related Commands	show ib sm routing-engines	
Notes	The list of root GIDs are relevant when IB SM is running on the switch, and the routing algorithm is up-down or fat-tree.	

10.5.6.1.44 ib sm routing-engines

	ib sm routing-engines [dor] [file] [ftree] [lash] [minhop] [none] [updn] [ar-updn] no ib sm routing-engines Sets the routing engine of the SM. The no form of the command sets the routing engine to be “none”. The default SM routing engine is used.	
Syntax Description	dor	Includes “dor” engine in selection of routing engines
	file	Includes “file” engine in selection of routing engines
	ftree	Includes “ftree” engine in selection of routing engines

	lash	Includes “lash” engine in selection of routing engines
	minhop	Includes “minhop” engine in selection of routing engines
	none	No routing engines specified; use SM default(s)
	updn	Includes “up/down” engine in selection of routing engines
	ar-updn	Includes “adaptive routing up/down” engine in selection of routing engines
Default	None	
Configuration Mode	config	
History	3.1.0000	
	3.10.4000	Added ar-updn option
Example	switch (config) # ib sm routing-engines none	
Related Commands	show ib sm routing-engines	
Notes	Multiple routing engines can be specified separated by spaces so that specific ordering of routing algorithms will be tried if earlier routing engines fail.	

10.5.6.1.45 ib sm rtr-aguid-enable

	ib sm rtr-aguid-enable <value> no ib sm rtr-aguid-enable Configures SM alias GUID control option. The no form of the command resets SM alias GUID control to its default value.	
Syntax Description	value	Possible values: <ul style="list-style-type: none"> • 0—does not configure alias GIDs required by routers • 1—configures alias GIDs required by routers • 2—clears and does not configure alias GIDs required by routers
Default	0	
Configuration Mode	config	
History	3.6.2002	
Example	switch (config) # ib sm rtr-aguid-enable 1	
Related Commands		
Notes		

10.5.6.1.46 ib sm rtr-pr-flow-label

	ib sm rtr-pr-flow-label <value> no ib sm rtr-pr-flow-label <value> Configures inter-subnet PathRecord FlowLabel. The no form of the command resets inter-subnet PathRecord FlowLabel to its default value.	
Syntax Description	value	Range: 0-1048575
Default	0	
Configuration Mode	config	

History	3.6.2002
Example	switch (config) # ib sm rtr-pr-flow-label 1
Related Commands	
Notes	

10.5.6.1.47 ib sm rtr-pr-mtu

	ib sm rtr-pr-mtu <value> no ib sm rtr-pr-mtu <value> Configures inter-subnet PathRecord MTU. The no form of the command resets inter-subnet PathRecord MTU to its default value.	
Syntax Description	value	Possible values: 256, 512, 1K, 2K, 4K
Default	2K	
Configuration Mode	config	
History	3.6.2002	
Example	switch (config) # ib sm rtr-pr-mtu 2k	
Related Commands		
Notes		

10.5.6.1.48 ib sm rtr-pr-rate

	ib sm rtr-pr-rate <value> no ib sm rtr-pr-rate <value> Configures inter-subnet PathRecord rate. The no form of the command resets inter-subnet PathRecord rate to its default value.	
Syntax Description	value	Possible values: 2.5, 5, 10, 14, 20, 25, 40, 56, 100
Default	100	
Configuration Mode	config	
History	3.6.2002	
Example	switch (config) # ib sm rtr-pr-rate 5	
Related Commands		
Notes		

10.5.6.1.49 ib sm rtr-pr-sl

	ib sm rtr-pr-sl <value> no ib sm rtr-pr-sl <value> Configures inter-subnet PathRecord SL. The no form of the command resets inter-subnet PathRecord SL to its default value.	
Syntax Description	value	Range: [0-15]
Default	0	

Configuration Mode	config
History	3.6.2002
Example	switch (config) # rtr-pr-sl 0
Related Commands	
Notes	

10.5.6.1.50 ib sm rtr-pr-tclass

	ib sm rtr-pr-tclass <value> no ib sm rtr-pr-tclass <value> Configures inter-subnet PathRecord T-class. The no form of the command resets inter-subnet PathRecord T-class to its default value.	
Syntax Description	value	Range: 0-255
Default	0	
Configuration Mode	config	
History	3.6.2002	
Example	switch (config) # ib sm rtr-pr-tclass 1	
Related Commands		
Notes		

10.5.6.1.51 ib sm sa-key

	ib sm sa-key <SA_Key> Sets the SA_Key 64-bit value used by SA to qualify that a query is “trusted”.	
Syntax Description	SA Key	64 bit
Default	00:00:00:00:00:00:00:01	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm sa-key 5	
Related Commands	show ib sm sa-key	
Notes	OpenSM version 3.2.1 and lower used the default value of “1” in host byte order. You may need to change this value to inter-operate with older subnet managers.	

10.5.6.1.52 ib sm single-thread

	ib sm single-thread no ib sm single-thread Enables the Subnet Manager to use a single thread to service all requests. The no form of the command enables SA to use multiple service threads.	
Syntax Description	N/A	
Default	Disabled (use multiple service threads)	

Configuration Mode	config
History	3.1.0000
Example	switch (config) # ib sm single-thread
Related Commands	show ib sm single-thread
Notes	

10.5.6.1.53 ib sm sm-inactive

	ib sm sm-inactive no ib sm sm-inactive Configures the SM to start in the “inactive” SM state. This option can be used to run a standalone system without the SM/SA function. The no form of the command configures the SM to start in “init” SM state.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm sm-inactive	
Related Commands	show ib sm sm-inactive	
Notes		

10.5.6.1.54 ib sm sm-key

	ib sm sm-key <SM_Key> Sets the SM 64-bit SM_Key.	
Syntax Description	SM Key	64 bit
Default	00:00:00:00:00:00:00:01	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm sm-key 00:00:00:00:00:00:00:05	
Related Commands	show ib sm sm-key	
Notes	OpenSM version 3.2.1 and lower used the default value of “1” in host byte order. You may need to change this value to inter-operate with older subnet managers.	

10.5.6.1.55 ib sm sm-priority

	ib sm sm-priority <priority> Prioritizes the desired SM compared to other SMs on the fabric.	
Syntax Description	priority	Range: 0-15 0 is least important 15 the most important
Default	0	

Configuration Mode	config
History	3.1.0000
Example	switch (config) # ib sm sm-priority 1
Related Commands	show ib sm sm-priority
Notes	If two or more active SMs have the same highest priority, the one with the lowest port GUID manages the fabric.

10.5.6.1.56 ib sm sm-sl

	ib sm sm-sl <sm-sl> Sets the SM service level for SM/SA communication.	
Syntax Description	sm-sl	0-15
Default	0	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm sm-sl 10	
Related Commands	show ib sm sm-sl	
Notes	Selects the SL that is used for MADs.	

10.5.6.1.57 ib sm sminfo-poll-time

	ib sm sminfo-poll-time <time> This variable controls the timeout between two polls of an active subnet manager.	
Syntax Description	time	In milliseconds
Default	10 seconds	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm sminfo-poll-time 15	
Related Commands	show ib sm sminfo-poll-time	
Notes		

10.5.6.1.58 ib sm subnet-prefix

	ib sm subnet-prefix <prefix> no ib sm subnet-prefix <prefix> Sets the SM "Subnet Prefix" used to create scope qualifiers for all elements managed by the SM. The no form of the command resets the subnet prefix to its default value.	
Syntax Description	prefix	64 bit
Default	FE:80:00:00:00:00:00	
Configuration Mode	config	

History	3.6.1002	
	3.6.2002	Added no form of the command
Example	<code>switch (config) # ib sm subnet-prefix ff:ff:ff:ff:ff:ff:ff:00</code>	
Related Commands	show ib sm subnet-prefix	
Notes	The default value is also the InfiniBand default for a locally administered subnet.	

10.5.6.1.59 ib sm subnet-prefix-override

	ib sm subnet-prefix-override no ib sm subnet-prefix-override Disables IB Router subnet prefix checking. The no form of the command enables IB Router subnet prefix checking.	
Syntax Description	N/A	
Default	Enabled	
Configuration Mode	config	
History	3.6.2002	
Example	<code>switch (config) # ib sm subnet-prefix-override</code>	
Related Commands	show ib sm subnet-prefix-override	
Notes		

10.5.6.1.60 ib sm max-smpps-timeout

	ib sm max-smpps-timeout <Timeout> Sets timeout for SMPs between max_wire_smpps & max_wire_smpps2	
Syntax Description	timeout	Timeout in seconds
Default	N/A	
Configuration Mode	config	
History	3.8.3000	
Example	<code>switch (config) # ib sm max-smpps-timeout 22</code>	
Related Commands		
Notes		

10.5.6.1.61 ib sm subnet-timeout

	ib sm subnet-timeout <time> Sets the global per-port subnet timeout value (PortInfo:SubnetTimeOut). This value also controls the maximum trap frequency in which no traps are allowed to be sent faster than the subnet_timeout value.	
Syntax Description	time	The actual timeout is $4.096 \mu\text{s} * 2^{*}\langle\text{time}\rangle$. The range of time is 0-31 for this parameter which supports 32 discrete time values between 4 μs and about 2.4 hours.
Default	0x12 (About 1 second)	

Configuration Mode	config
History	3.1.0000
Example	<code>switch (config) # ib sm subnet-timeout 5</code>
Related Commands	show ib sm subnet-timeout
Notes	If the SMA generates a sequence of traps, the interval between successive traps should not be smaller than <time>.

10.5.6.1.62 ib sm sweep-interval

	ib sm sweep-interval <time> no ib sm sweep-interval Specifies the time between subnet sweeps. The no form of the command disables periodic sweeps.	
Syntax Description	time	Range: Between 0 and 36000 seconds; 0—disable
Default	10 seconds	
Configuration Mode	config	
History	3.1.0000	
Example	<code>switch (config) # ib sm sweep-interval 20</code>	
Related Commands	show ib sm sweep-interval	
Notes		

10.5.6.1.63 ib sm sweep-on-trap

	ib sm sweep-on-trap no ib sm sweep-on-trap Enables every TRAP received by the SM to initiate a heavy sweep in addition to the processing required by the TRAP. The no form of the command enables SM to use a combination of light and heavy sweeps based on the type of TRAP and other internal states.	
Syntax Description	N/A	
Default	enable	
Configuration Mode	config	
History	3.1.0000	
Example	<code>switch (config) # ib sm sweep-on-trap</code>	
Related Commands	show ib sm sweep-on-trap	
Notes	More than 10 successive identical TRAPs disable the automatic sweep behavior until at least one different TRAP has been received.	

10.5.6.1.64 ib sm transaction-retries

	ib sm transaction-retries <transaction-retries-count> Sets the maximum retries for failed transactions.	
--	--	--

Syntax Description	transaction-retries-count	Must be an integer
Default	3	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm transaction-retries 10	
Related Commands	show ib sm transaction-retries	
Notes		

10.5.6.1.65 ib sm use-heavy-sweeps

	ib sm use-heavy-sweeps no ib sm use-heavy-sweeps Turns every fabric sweep to a heavy sweep. The no form of the command enables the SM to use a combination of light and heavy sweeps.	
Syntax Description	N/A	
Default	disable	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm use-heavy-sweeps	
Related Commands	show ib sm use-heavy-sweeps	
Notes		

10.5.6.1.66 ib sm use-ucast-cache

	ib sm use-ucast-cache no ib sm use-ucast-cache Enables the SM to use cached routine data (LMC=0 only). The no form of the command disables the SM to use cached routine data.	
Syntax Description	N/A	
Default	Disabled	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib sm use-ucast-cache	
Related Commands	show ib sm use-ucast-cache	
Notes		

10.5.6.1.67 `ib sm vl-stalls`

	<code>ib sm vl-stalls <count></code> Sets the number of sequential frame drops that cause a switch-to-switch port to enter the VLStalled state.	
Syntax Description	count	1-255
Default	7	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # ib sm vl-stalls 10</pre>	
Related Commands	show ib sm vl-stalls	
Notes		

10.5.6.1.68 `ib sm virt`

	<code>ib sm virt {enable disable ignore}</code> <code>no ib sm virt</code> Configures IB SM port virtualization support. The no form of the command resets this parameter to its default value.	
Syntax Description	enable	IB SM supports virtualization, and configures virtual ports
	disable	IB SM disables virtual ports
	ignore	IB SM ignores virtual ports and does not change their configuration
Default	ignore	
Configuration Mode	config	
History	3.4.2008	
Example	<pre>switch (config) # ib sm virt configure</pre>	
Related Commands		
Notes		

10.5.6.1.69 `ib sm virt-default-hop-limit`

	<code>ib sm virt-default-hop-limit <value></code> <code>no ib sm virt-default-hop-limit</code> Configures the default value for hop limit to be returned in path records. The no form of the command resets this parameter to its default value.	
Syntax Description	value	Range: 0-255
Default	2	
Configuration Mode	config	
History	3.6.2002	
Example	<pre>switch (config) # ib sm virt-default-hop-limit 3</pre>	
Related Commands		

Notes	
-------	--

10.5.6.1.70 ib sm virt-max-ports-in-process

	ib sm virt-max-ports-in-process <value> no ib sm virt-max-ports-in-process Configures the maximum number of ports to be processed simultaneously. The no form of the command resets this parameter to its default value.	
Syntax Description	value	Range: 0-65535 "0" processes all pending ports
Default	4	
Configuration Mode	config	
History	3.6.2002	
Example	switch (config) # ib sm virt-max-ports-in-process 5	
Related Commands		
Notes		

10.5.6.2 Show

10.5.6.2.1 show ib sm

	show ib sm Displays the SM admin state.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example	switch (config) # show ib sm enable	
Related Commands	ib sm	
Notes		

10.5.6.2.2 show ib sm accum-log-file

	show ib sm accum-log-file Displays the accum-log-file configuration.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example	switch (config) # show ib sm accum-log-file enable	

Related Commands	ib sm accum-log-file
Notes	

10.5.6.2.3 show ib sm babbling-policy

	show ib sm babbling-policy Displays the ability of the SM to disable babbling ports (i.e., generating frequent traps).
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config) # show ib sm babbling-policy disable
Related Commands	ib sm babbling-policy
Notes	

10.5.6.2.4 show ib sm calculate-missing-routes

	Show ib sm calculate-missing-routes Display option allowing SM to find and recalculate missing routes without creating credit-loops
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.8.2000
Example	switch (config) # ib sm calculate-missing-routes switch (config) # show ib sm calculate-missing-routes ib sm calculate-missing-routes: enabled switch (config) # no ib sm calculate-missing-routes switch (config) # show ib sm calculate-missing-routes ib sm calculate-missing-routes: disabled
Related Commands	ib sm calculate-missing-routes
Notes	

10.5.6.2.5 show ib sm connect-roots

	show ib sm connect-roots Displays the IBA compliant multi-stage switch directive.
Syntax Description	N/A
Default	N/A

Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config) # show ib sm connect-roots true
Related Commands	ib sm connect-roots
Notes	

10.5.6.2.6 show ib sm enable-quirks

	show ib sm enable-quirks Displays if the SM uses high risk features and handles HW workarounds.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config) # show ib sm enable-quirks disable
Related Commands	ib sm enable-quirks
Notes	

10.5.6.2.7 show ib sm exit-on-fatal

	show ib sm exit-on-fatal Displays if the SM exits upon a fatal error.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config) # show ib sm exit-on-fatal enable
Related Commands	ib sm exit-on-fatal
Notes	

10.5.6.2.8 show ib sm fdr10

	show ib sm fdr10 Displays the status of the SM use of FDR10.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode

History	3.1.0000
Example	switch (config) # show ib sm fdr10 SM use of fdr10 is off
Related Commands	
Notes	

10.5.6.2.9 show ib sm force-link-speed

	show ib sm force-link-speed Displays SM behavior for PortInfo:LinkSpeedEnabled parameter on switch ports.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.4.1604	Updated Syntax Description, Example and Notes
Example	switch (config) # show ib sm force-link-speed Default: set to PortInfo:LinkSpeedSupported	
Related Commands	ib sm force-link-speed	
Notes	Possible outputs: <ul style="list-style-type: none"> • Default: set to PortInfo:LinkSpeedExtSupported • Disabled: extended link speed not in use • Negotiate: <a list containing fdr, edr speeds> 	

10.5.6.2.10 show ib sm force-link-speed-ext

	show ib sm force-link-speed-ext Displays SM behavior for PortInfo:LinkSpeedExtEnabled parameter on the switch ports.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.4.1604	Updated Description and Example
Example	switch (config) # show ib sm force-link-speed-ext Negotiate: fdr edr	
Related Commands	ib sm force-link-speed	
Notes	Possible outputs: <ul style="list-style-type: none"> • Default: set to PortInfo:LinkSpeedExtSupported • Disabled: extended link speed not in use • Negotiate: <a list containing fdr, edr speeds> 	

10.5.6.2.11 show ib sm force-log-flush

	show ib sm force-log-flush Displays if every log message generated forces the log to be flushed.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
	3.4.1604	Updated Description and Example
Example	<pre>switch (config) # show ib sm force-log-flush enable</pre>	
Related Commands	ib sm force-log-flush	
Notes		

10.5.6.2.12 show ib sm guid2lid-cache

	show ib sm guid2lid-cache Displays whether or not the SM honors the cached GUID-to-LID mapping information.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example	<pre>switch (config) # show ib sm guid2lid-cache disable</pre>	
Related Commands	ib sm guid2-lid-cache	
Notes		

10.5.6.2.13 show ib sm honor-partitions

	show ib sm honor-partitions Displays the partition enforcement settings in the subnet.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example	<pre>switch (config) # show ib sm honor-partitions disable</pre>	
Related Commands	ib sm honor-partitions	
Notes		

10.5.6.2.14 show ib sm hoq-lifetime

	<code>show ib sm hoq-lifetime</code> Displays the maximum time a frame can wait at the head of a switch-to-switch port queue before it is dropped.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm hoq-lifetime 0x12 (About 1 second)</pre>
Related Commands	<code>ib sm hoq-lifetime</code>
Notes	

10.5.6.2.15 show ib sm ignore-other-sm

	<code>show ib sm ignore-other-sm</code> Displays if the rules governing SM elections and attempt to manage the fabric on the node are ignored by the SM.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm ignore-other-sm enable</pre>
Related Commands	<code>ib sm ignore-other-sm</code>
Notes	

10.5.6.2.16 show ib sm ipv6-nsm

	<code>show ib sm ipv6-nsm</code> Displays the consolidation of IPv6 Solicited Node Multicast (SNM) group join requests.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm ipv6-nsm enable</pre>
Related Commands	<code>ib sm ipv6-nsm</code>
Notes	

10.5.6.2.17 show ib sm lash

	<code>show ib sm lash {do-mesh-analysis start-vl}</code> Display “lash” routing method parameters.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm lash do-mesh-analysis enable</pre>
Related Commands	<code>ib sm lash</code>
Notes	

10.5.6.2.18 show ib sm leafhoq-lifetime

	<code>show ib sm leafhoq-lifetime</code> Displays the maximum time a frame can wait at the head of a switch-to-CA_or_Router port queue before it is dropped.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm leafhoq-lifetime 0x10 (About 268 mS)</pre>
Related Commands	<code>ib sm leafhoq-lifetime</code>
Notes	

10.5.6.2.19 show ib sm leafvl-stalls

	<code>show ib sm leafvl-stalls</code> Displays the number of sequential frame drops that case a switch-to-CA_or_Router port to enter the VLStalled state.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm leafvl-stalls 7</pre>
Related Commands	<code>ib sm leafvl-stalls</code>
Notes	

10.5.6.2.20 show ib sm lmc

	show ib sm lmc Displays the number of sequential frame drops that case a switch-to-CA_or_Router port to enter the VLStalled state.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config) # show ib sm lmc 0x0
Related Commands	ib sm lmc
Notes	

10.5.6.2.21 show ib sm lmc-esp0

	show ib sm lmc-esp0 Displays the number of sequential frame drops that case a switch-to-CA_or_Router port to enter the VLStalled state.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config) # show ib sm lmc-esp0 enable
Related Commands	ib sm lmc-esp0
Notes	

10.5.6.2.22 show ib sm log

	show ib sm log [continuous] [[not] [matching <reg-expression>]] Displays IB SM event logs.	
Syntax Description	continuous	Displays IB SM new event log messages as they arrive
	not	Displays IB SM new event logs that do not match a given regular expression
	matching	Displays IB SM event log messages that match a given regular expression
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	

Example	<pre>switch (config) # show ib sm log Jul 18 12:00:40 165863 [48026660] 0x03 -> OpenSM 3.3.13.MLNX_20121224_9b362db Jul 18 12:00:40 168685 [48026660] 0x80 -> OpenSM 3.3.13.MLNX_20121224_9b362db Jul 18 12:00:40 170789 [48026660] 0x02 -> osm_vendor_init: 1000 pending umads specified Jul 18 12:00:40 175696 [48026660] 0x80 -> Entering DISCOVERING state Jul 18 12:00:40 249448 [48026660] 0x02 -> osm_vendor_bind: Binding to port 0x2c903008b0440 Jul 18 12:00:40 293959 [48026660] 0x02 -> osm_vendor_bind: Binding to port 0x2c903008b0440 Jul 18 12:00:40 296921 [48026660] 0x02 -> osm_vendor_bind: Binding to port 0x2c903008b0440 Jul 18 12:00:40 304702 [48026660] 0x02 -> osm_opensm_bind: Setting IS_SM on port 0x0002c903008b0440 Jul 18 12:00:40 399744 [4A85D4B0] 0x80 -> Entering MASTER state</pre>
Related Commands	show ib sm log-flags
Notes	

10.5.6.2.23 show ib sm log-flags

	show ib sm log-flags Displays what type of messages the SM is logging.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm log-flags 0x3 (error, info)</pre>
Related Commands	ib sm log-flags
Notes	

10.5.6.2.24 show ib sm log-max-size

	show ib sm log-max-size Displays the maximum size of the log file.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Example	<pre>switch (config) # show ib sm log-max-size 50 MBytes</pre>
Related Commands	ib sm log-max-size
Notes	

10.5.6.2.25 show ib sm max-op-vls

	<code>show ib sm max-op-vls</code> Displays the maximum size of the log file.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm max-op-vls 15</pre>
Related Commands	<code>ib sm max-op-vls</code>
Notes	

10.5.6.2.26 show ib sm max-ports

	<code>show ib sm max-ports</code> Displays the number of CA ports SM can manage.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm max-ports 2048</pre>
Related Commands	<code>ib sm max-ports</code>
Notes	

10.5.6.2.27 show ib sm max-reply-time

	<code>show ib sm max-reply-time</code> Displays the number of CA ports SM can manage.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm max-reply-time 200 milliseconds</pre>
Related Commands	<code>ib sm max-reply-time</code>
Notes	

10.5.6.2.28 show ib sm max-reverse-hops

	<code>show ib sm max-reverse-hops</code> Displays max hops IO node to top switch.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm max-reverse-hops 0</pre>
Related Commands	<code>ib sm max-reverse-hops</code>
Notes	

10.5.6.2.29 show ib sm aguid-default-hop-limit

	<code>show ib sm aguid-default-hop-limit</code> Displays max hops IO node to top switch.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm aguid-default-hop-limit 1</pre>
Related Commands	<code>ib sm aguid-default-hop-limit</code>
Notes	

10.5.6.2.30 show ib sm max-wire-smpls

	<code>show ib sm max-wire-smpls</code> Displays the maximal number of MADs the SM will have outstanding at one time to count.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm max-wire-smpls 4</pre>
Related Commands	<code>ib sm max-wire-smpls</code>
Notes	

10.5.6.2.31 show ib sm max-wire-smpls2

	show ib sm max-wire-smpls2 Displays maximal SM timeout based packets allowed to be outstanding.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config) # show ib sm max-wire-smpls2 4
Related Commands	ib sm max-wire-smpls2
Notes	

10.5.6.2.32 show ib sm mkey-lease

	show ib sm mkey-lease Displays MKey period in seconds.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config) # show ib sm mkey-lease 0 (no timeout)
Related Commands	ib sm mkey-lease
Notes	

10.5.6.2.33 show ib sm m-key

	show ib sm m-key Displays the MKey used by the SM.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
	3.6.2002 Updated Example
Example	switch (config) # show ib sm m-key 11:33:55:77:99:aa:cc:ee
Related Commands	ib sm m-key
Notes	

10.5.6.2.34 show ib sm mkey-lease

	<code>show ib sm mkey-lease</code> Displays MKey lease period in seconds.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.2002
Example	<pre>switch (config) # show ib sm mkey-lease 0 (No timeout)</pre>
Related Commands	<code>ib sm mkey-lookup</code>
Notes	

10.5.6.2.35 show ib sm mkey-lookup

	<code>show ib sm mkey-lookup</code> Displays whether the SM looks in file cache for unknown note MKeys.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.2002
Example	<pre>switch (config) # show ib sm mkey-lookup enable</pre>
Related Commands	<code>ib sm m-key</code>
Notes	

10.5.6.2.36 show ib sm mkey-protect-level

	<code>show ib sm mkey-protect-level</code> Displays MKey protection level.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.2002
Example	<pre>switch (config) # show ib sm mkey-protect-level 0</pre>
Related Commands	<code>ib sm mkey-protect-level</code>
Notes	

10.5.6.2.37 show ib sm msgfifo-timeout

	show ib sm msgfifo-timeout Displays the elapsed time in milliseconds before a frame at the head of Subnet Agent queue causes an immediate BUSY state.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm msgfifo-timeout 10.000 seconds</pre>
Related Commands	ib sm msgfifo-timeout
Notes	

10.5.6.2.38 show ib sm multicast

	show ib sm multicast Displays whether the SM supports multicast on the fabric.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm multicast enable</pre>
Related Commands	ib sm multicast
Notes	

10.5.6.2.39 show ib sm no-client-rereg

	show ib sm no-client-rereg Displays client re-registration admin state.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm no-client-rereg enable</pre>
Related Commands	ib sm no-client-rereg
Notes	

10.5.6.2.40 show ib sm overrun-trigger

	<code>show ib sm overrun-trigger</code> Displays count of local buffer overrun errors for Infiniband trap 130.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm overrun-trigger 3</pre>
Related Commands	<code>ib sm overrun-trigger</code>
Notes	

10.5.6.2.41 show ib sm packet-life-time

	<code>show ib sm packet-life-time</code> Displays the maximum time a frame can live in a switch.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm packet-life-time 0x14 (Infinite)</pre>
Related Commands	<code>ib sm packet-life-time</code>
Notes	

10.5.6.2.42 show ib sm phy-err-trigger

	<code>show ib sm phy-err-trigger</code> Displays the number of local link integrity errors and the port's SMA supports traps.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm phy-err-trigger 5</pre>
Related Commands	<code>ib sm phy-err-trigger</code>
Notes	

10.5.6.2.43 show ib sm polling-retries

	<code>show ib sm polling-retries</code> Displays the number of consecutive times an active SM must fail to respond before it is declared dead.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm polling-retries 8</pre>
Related Commands	<code>ib sm polling-retries</code>
Notes	

10.5.6.2.44 show ib sm port-prof-switch

	<code>show ib sm port-prof-switch</code> Displays whether or not the counting of adapters, routers, and switches through the links is being done.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm port-prof-switch true</pre>
Related Commands	<code>ib sm port-prof-switch</code>
Notes	

10.5.6.2.45 show ib sm reassign-lids

	<code>show ib sm reassign-lids</code> Displays the ability of the SM to reassign LIDs to nodes it finds already configured with a valid LID.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config) # show ib sm reassign-lids enable</pre>
Related Commands	<code>ib sm reassign-lids</code>
Notes	

10.5.6.2.46 show ib sm root-guid

	<code>show ib sm root-guid</code> Displays the configured root GUIDs for the SM.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config)# show ib sm root-guid AA:00:11:22:33:44:55 AA:00:11:22:33:44:56 AA:00:11:22:33:44:57 ...</pre>
Related Commands	<code>ib sm routing-engine</code>
Notes	The list of root GUIDs are relevant when IB SM is running on the switch, and the routing algorithm is up-down or fat-tree.

10.5.6.2.47 show ib sm routing-engines

	<code>show ib sm routing-engines</code> Displays an ordered list of routing engines.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config)# show ib sm routing-engines none</pre>
Related Commands	<code>ib sm routing-engines</code>
Notes	

10.5.6.2.48 show ib sm routing-info

	<code>show ib sm routing-info</code> Displays current routing engine information.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config)# show ib sm routing-info Current routing engine minhop</pre>
Related Commands	
Notes	

10.5.6.2.49 show ib sm rtr-aguid-enable

	show ib sm rtr-aguid-enable Displays GUID option configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.2002
Example	switch (config)# show ib sm rtr-aguid-enable 0
Related Commands	ib sm rtr-aguid-enable
Notes	

10.5.6.2.50 show ib sm rtr-pr-flow-label

	show ib sm rtr-pr-flow-label Displays inter-subnet PathRecord FlowLabel.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.2002
Example	switch (config)# show ib sm rtr-pr-flow-label 0
Related Commands	ib sm rtr-pr-flow-label
Notes	"0" signifies that inter-subnet PathRecord FlowLabel is disabled

10.5.6.2.51 show ib sm rtr-pr-mtu

	show ib sm rtr-pr-mtu Displays inter-subnet PathRecord MTU.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.2002
Example	switch (config)# show ib sm rtr-pr-mtu 2K
Related Commands	ib sm rtr-pr-mtu
Notes	

10.5.6.2.52 show ib sm rtr-pr-rate

	<code>show ib sm rtr-pr-rate</code> Displays inter-subnet PR rate.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.2002
Example	<pre>switch (config)# show ib sm rtr-pr-rate 100</pre>
Related Commands	<code>ib sm rtr-pr-rate</code>
Notes	

10.5.6.2.53 show ib sm rtr-pr-sl

	<code>show ib sm rtr-pr-sl</code> Displays inter-subnet PathRecord service level.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.2002
Example	<pre>switch (config)# show ib sm rtr-pr-sl 0</pre>
Related Commands	<code>ib sm rtr-pr-sl</code>
Notes	

10.5.6.2.54 show ib sm sa-key

	<code>show ib sm sa-key</code> Displays the SM sa-key value used by SA to qualify that a query is “trusted”.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config)# show ib sm sa-key 00:00:00:00:00:00:00:05</pre>
Related Commands	<code>ib sm rtr-pr-sl</code>
Notes	

10.5.6.2.55 show ib sm single-thread

	<code>show ib sm single-thread</code> Displays if the SM uses a single thread to service all requests.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<code>switch (config)# show ib sm single-thread enable</code>
Related Commands	<code>ib sm single-thread</code>
Notes	

10.5.6.2.56 show ib sm sm-inactive

	<code>show ib sm sm-inactive</code> Displays whether or not the SM starts in “inactive” rather than “init” SM state.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<code>switch (config)# show ib sm sm-inactive enable</code>
Related Commands	<code>ib sm sm-inactive</code>
Notes	

10.5.6.2.57 show ib sm sm-key

	<code>show ib sm sm-key</code> Displays the SM 64-bit SM_Key.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.63.1.0000
Example	<code>switch (config)# show ib sm sm-key 00:00:00:00:00:00:00:05</code>
Related Commands	<code>ib sm sm-key</code>
Notes	

10.5.6.2.58 show ib sm sm-priority

	<code>show ib sm sm-priority</code> Displays the importance of this SM compared to other SMs on the fabric.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config)# show ib sm sm-priority 1</pre>
Related Commands	<code>ib sm sm-priority</code>
Notes	If 2 or more active SMs have the same highest priority, the one with the lowest port GUID will manage the fabric.

10.5.6.2.59 show ib sm sm-sl

	<code>show ib sm sm-sl</code> Displays SL used for SM/SA communication.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config)# show ib sm sm-sl 1</pre>
Related Commands	<code>ib sm sm-sl</code>
Notes	

10.5.6.2.60 show ib sm sminfo-poll-time

	<code>show ib sm sminfo-poll-time</code> Displays the timeout in milliseconds between two polls of an active SM.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config)# show ib sm sminfo-poll-time 15 milliseconds</pre>
Related Commands	<code>ib sm sminfo-poll-time</code>
Notes	

10.5.6.2.61 show ib sm subnet-prefix

	show ib sm subnet-prefix Displays the SM “Subnet Prefix” used to create scope qualifiers for all elements managed by the SM.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config)# show ib sm subnet-prefix FF:FF:FF:FF:FF:FF:FF:00
Related Commands	ib sm subnet-prefix
Notes	

10.5.6.2.62 show ib sm subnet-prefix-override

	show ib sm subnet-prefix-override Displays whether IB Router subnet prefix checking is enabled or disabled.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config)# show ib sm subnet-prefix-override disable
Related Commands	ib sm subnet-prefix-override
Notes	

10.5.6.2.63 show ib sm subnet-timeout

	show ib sm subnet-timeout Displays the global per-port subnet timeout value (PortInfo:SubnetTimeOut). This value also controls the maximum trap frequency in which no traps are allowed to be sent faster than the subnet_timeout value. The time is 4.096 uS * 2*time.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	switch (config)# show ib sm subnet-timeout 0x5 (About 131 uS)
Related Commands	ib sm subnet-timeout
Notes	

10.5.6.2.64 show ib sm sweep-interval

	<code>show ib sm sweep-interval</code> Displays the time in seconds between subnet sweeps.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config)# show ib sm sweep-interval 20 seconds</pre>
Related Commands	<code>ib sm sweep-interval</code>
Notes	

10.5.6.2.65 show ib sm sweep-on-trap

	<code>show ib sm sweep-on-trap</code> Displays whether or not a heavy sweep is initiated by the TRAP received by the SM.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config)# show ib sm sweep-on-trap enable</pre>
Related Commands	<code>ib sm sweep-on-trap</code>
Notes	

10.5.6.2.66 show ib sm transaction-retries

	<code>show ib sm transaction-retries</code> Displays maximum retries before failing a transaction.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config)# show ib sm transaction-retries 3</pre>
Related Commands	<code>ib sm transaction-retries</code>
Notes	

10.5.6.2.67 show ib sm use-heavy-sweeps

	<code>show ib sm use-heavy-sweeps</code> Displays maximum retries before failing a transaction.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config)# show ib sm use-heavy-sweeps disable</pre>
Related Commands	<code>ib sm use-heavy-sweeps</code>
Notes	

10.5.6.2.68 show ib sm use-ucast-cache

	<code>show ib sm use-ucast-cache</code> Displays maximum retries before failing a transaction.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config)# show ib sm use-ucast-cache false</pre>
Related Commands	<code>ib sm use-ucast-cache</code>
Notes	

10.5.6.2.69 show ib sm version

	<code>show ib sm version</code> Displays the OpenSM version currently running.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.2.3000
Example	<pre>switch (config)# show ib sm version OpenSM5.2.0</pre>
Related Commands	
Notes	

10.5.6.2.70 show ib sm virt-default-hop-limit

	<code>show ib sm virt-default-hop-limit</code> Displays the open SM version that is currently running.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.2002
Example	<pre>switch (config)# show ib sm virt-default-hop-limit 2</pre>
Related Commands	<code>ib sm virt-default-hop-limit</code>
Notes	

10.5.6.2.71 show ib sm virt-max-ports-in-process

	<code>show ib sm virt-max-ports-in-process</code> Displays the maximum number of ports to be processed simultaneously.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.2002
Example	<pre>switch (config)# show ib sm virt-max-ports-in-process 4</pre>
Related Commands	<code>ib sm virt-max-ports-in-process</code>
Notes	

10.5.6.2.72 show ib sm vl-stalls

	<code>show ib sm use-vl-stalls</code> Displays the number of sequential frame drops that cause a switch-to-switch port to enter the VLStalled state.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000
Example	<pre>switch (config)# show ib sm vl-stalls 7</pre>
Related Commands	<code>ib sm vl-stalls</code>
Notes	

10.5.6.3 Partitions

10.5.6.3.1 ib partition

	ib partition <partition-name> [pkey <pkey number>] no ib partition <partition-name> [force] Enters the context of the partition specified. The no form of the command deletes the partition.	
Syntax Description	partition-name	Name of partition context to be entered
	pkey	Creates a partition and enters a new configuration mode
	force	Forces configuration
Default	Default partition is available (PKEY 0x7fff)	
Configuration Mode	config	
History	3.2.0500	
	3.6.8008	
Example	<pre>switch (config) # ib partition my-partition switch (config partition my-partition) #</pre>	
Related Commands		
Notes		

10.5.6.3.2 pkey

	pkey <number> [force] no pkey <number> Specifies PKEY number for this partition. The no form of the command removes the PKEY configuration from partitions.conf file.	
Syntax Description	number	Range: 0x001-0x7fff
	force	Forces configuration
Default	N/A	
Configuration Mode	config partition	
History	3.2.0500	
	3.5.1000	Added “force” parameter
Example	<pre>switch (config partition my-partition) # pkey 0x7777</pre>	
Related Commands		
Notes	PKEY must be unique	

10.5.6.3.3 defmember

	defmember <type> [force] no defmember Sets the default membership for port GUID list. The no form of the command set the defmember configuration to default (it will not appear in the partitions.conf file).
--	--

Syntax Description	type	Default membership for GUIDs in this partition: <ul style="list-style-type: none"> • full • limited • both
	force	Forces configuration
Default	limited	
Configuration Mode	config partition	
History	3.2.0500	
	3.4.1100	Added “both” option
	3.5.1000	Added “force” parameter
Example	<code>switch (config ib partition my-partition) # defmember full</code>	
Related Commands	ib sm allow-both-pkeys member	
Notes	This parameter can be overwritten for specific GUID, using the “member” command.	

10.5.6.3.4 member

	<code>member {<guid> all all-cas all-routers all-switches all-vcas self } [type <member-type>] [force]</code> <code>no member {<guid> all all-cas all-routers all-switches all-vcas self } [type] [force]</code> Adds static members to partition. The no form of the command will remove the static member from the partition (it will not appear in the partitions.conf file).	
Syntax Description	guid	The GUID number
	all	Can be used for all GUIDs in the fabric
	self	Can be used for the the switch GUID
	all-cas	Adds all GUIDs that belong to CA ports in the fabric
	all-routers	Adds all GUIDS that belong to routers in the fabric
	all-switches	Adds all GUIDS that belong to switched in the fabric
	all-vcas	Adds all GUIDS that belong to virtual CA posts in the fabric
	member-type	Default membership for GUIDs in this partition: <ul style="list-style-type: none"> • full • limited • both
	force	Forces configuration (only relevant to the default partition)
Default	N/A	
Configuration Mode	config partition	
History	3.2.0500	
	3.4.1100	Added “both” parameter
	3.5.1000	Added “force” parameter
	3.8.2100	Added "all-cas," "all-routers," all-switches," and "all-vcas" parameters

Example	<code>switch (config ib partition my-partition) # member all</code>
Related Commands	<code>ib partition</code> <code>ib sm allow-both-pkeys</code> <code>defmember</code>
Notes	

10.5.6.3.5 ipoib

	<code>ipoib [force]</code> <code>no ipoib [force]</code> Enables this partition to use IPoIB. As a result IPoIB multicast group will be created. The no form of the command removes the use of IPoIB in this partition (it will not appear in the partitions.conf file).	
Syntax Description	<code>force</code>	Forces configuration
Default	no IPoIB	
Configuration Mode	config partition	
History	3.2.0500	
	3.5.1000	Added “force” parameter
	3.6.8008	Added “force” parameter to no form
Example	<code>switch (config ib partition my-partition) # ipoib</code>	
Related Commands	<code>ib partition</code> <code>rate</code> <code>mtu</code> <code>sl</code> <code>scope</code>	
Notes	The commands “rate”, “mtu”, “sl” and “scope” can be used only when the IPoIB parameter is enabled.	

10.5.6.3.6 mtu

	<code>mtu <256, 512, 1K, 2K,4K> [force]</code> <code>no mtu</code> Specifies MTU for this IPoIB multicast group. The no form of the command sets the mtu to default (it will not appear in the partitions.conf file).	
Syntax Description	<code>force</code>	Forces configuration
Default	2K	
Configuration Mode	config partition	
History	3.2.0500	
	3.5.1000	Added “force” parameter
Example	<code>switch (config ib partition my-partition) # mtu 4K</code>	
Related Commands	ipoib	
Notes	IPoIB parameter on the partitions must be enabled in order to use this parameter	

10.5.6.3.7 rate

	<code>rate <rate> [force]</code> <code>no rate</code> Specifies rate for this IPoB multicast group. The no form of the command set the rate to default (removes the rate from the partitions.conf).	
Syntax Description	<code>rate</code>	<ul style="list-style-type: none"> • default—Default • 2.5—2.5 Gbps • 5—5 Gbps • 10—10 Gbps • 14—14 Gbps • 20—20 Gbps • 25—25 Gbps • 40—40 Gbps • 56—56 Gbps • 100—100 Gbps
Default	10Gb/s	
Configuration Mode	config partition	
History	3.2.0500	
	3.4.1100	Updated rate Syntax Description
	3.5.1000	Added “force” parameter
Example	<pre>switch (config partition my-partition) # rate 20</pre>	
Related Commands	ipoib	
Notes	Ports that do not support the IPoB rate are not added to the partition	

10.5.6.3.8 scope

	<code>scope <type> [force]</code> <code>no scope <link-local, site-local, organization-local, global></code> Specifies scope for this IPoB multicast group. The no form of the command removes the scope configuration from the partitions.conf file.	
Syntax Description	<code>type</code>	<ul style="list-style-type: none"> • link-local • site-local • organization-local • global
	<code>force</code>	Forces configuration
Default	link-local	
Configuration Mode	config partition	
History	3.2.0500	
	3.5.1000	Added “force” parameter
Example	<pre>switch (config partition my-partition) # scope global</pre>	
Related Commands	ipoib	
Notes	IPoB parameter on the partitions must be enabled in order to use this parameter.	

10.5.6.3.9 sl

	sl <0-14, "default"> [force] no sl Specifies SL (Service Level - QoS) for this IPoB multicast group. The no form of the command sets it to default (the sl configuration is removed from the partitions.conf file).	
Syntax Description	force	Forces configuration
Default	Default (0)	
Configuration Mode	config partition	
History	3.2.0500	
	3.5.1000	Added "force" parameter
Example	<pre>switch (config partition my-partition) # sl 7</pre>	
Related Commands	ipoib	
Notes	IPoB parameter on the partitions must be enabled in order to use this parameter.	

10.5.6.3.10 show ib partition

	show ib partition [<partition-name> [member [<member-name>]]] Displays partition info, with optional to filters.	
Syntax Description	partition-name	Filters the output per partition name
	member <member-name>	Filters the output by a specific member
Default	N/A	
Configuration Mode	Any command mode	
History	3.2.0500	
	3.6.8008	Updated Example and note
Example	<pre>switch (config) # show ib partition Default Default PKey = 0x7FFF ipoib = yes members GUID='ALL' member='full'</pre>	
Related Commands		
Notes	If bulk update mode is enabled, this command notifies the user that these changes may not have been applied yet.	

10.5.6.4 Quality of Service (SM)

10.5.6.4.1 ib baseqos <port-type> high-limit

	ib baseqos <port-type> high-limit <count> Sets the high-limit value for the indicated port type. Thus the system will send at least 4096 * <count> bytes from the high priority list before sending any from the low priority list.
--	--

Syntax Description	port-type	<ul style="list-style-type: none"> ca—channel adapters rtr—routers sw0—ports 0 only of the switches swe—external ports of the switches
	high-limit	Possible values are: -1...255 <ul style="list-style-type: none"> -1—default SM high-limit 0—1 frame i =1...254 - 4K * i 255—unlimited
Default	-1 (default SM high-limit)	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib baseqos ca high-limit 255	
Related Commands	show ib baseqos	
Notes	A high-limit value of 255 means unlimited, and that makes it possible to starve the low priority list.	

10.5.6.4.2 ib baseqos max-vls

	ib baseqos <port-type> max-vls <value> Configures the maximum number of VLs for the indicated port type.	
Syntax Description	port-type	<ul style="list-style-type: none"> ca—channel adapters rtr—routers sw0—ports 0 only of the switches swe—external ports of the switches
	value	Range: 1-15
Default	15	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib baseqos ca max-vls 15	
Related Commands	show ib baseqos	
Notes		

10.5.6.4.3 ib baseqos sl2vl

	ib baseqos <port-type> sl2vl {sl0 sl0 sl1 sl0 sl1 sl2 ...} no ib baseqos <port-type> sl2vl Sets a list of up to 16 entries that map the SL entry to an appropriate VL. The no form of the command sets the attributes to their default settings.	
Syntax Description	port-type	<ul style="list-style-type: none"> ca—channel adapters rtr—routers sw0—ports 0 only of the switches swe—external ports of the switches

	sl[i]	A single vector (1 ... 16 elements), the command line vector determine the SL [0...15] that is mapped to the specified VL [0...15].
Default	The default mapping is: 0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,7	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # show ib baseqos ca sl2vl 0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,7 switch (config) # ib baseqos ca sl2vl 10 10 10 switch (config) # show ib baseqos ca sl2vl 10,10,10,15,15,15,15,15,15,15,15,15,15,15,15</pre>	
Related Commands	show ib baseqos	
Notes	Any missing SLs will be mapped to VL15.	

10.5.6.4.4 ib baseqos vlarb-high

	ib baseqos <port-type> vlarb-high {VW1 VW1 VW2 ...} no ib baseqos <port-type> vlarb-high Sets up to 15 VL to Weight mapping pairs for high priority processing. The no form of the command sets the attributes to their default settings.	
Syntax Description	port-type	<ul style="list-style-type: none"> ca—channel adapters rtr—routers sw0—ports 0 only of the switches swe—external ports of the switches
	VW[i]	There are two possible options for this parameter: <ul style="list-style-type: none"> A single vector (1 ...15) in the format of “#:#” separated by spaces, see example below. Format of “i#=X:Y” in order to change a specific entry (see example below)
Default	The default mapping is: 0:4,1:0,2:0,3:0,4:0,5:0,6:0,7:0,8:0,9:0,10:0,11:0,12:0,13:0,14:0	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # show ib baseqos ca vlarb-high 0:4,1:0,2:0,3:0,4:0,5:0,6:0,7:0,8:0,9:0,10:0,11:0,12:0,13:0,14:0 switch (config) # ib baseqos ca vlarb-high 0:10 1:10 switch (config) # show ib baseqos ca vlarb-high 0:10,1:10,2:0,3:0,4:0,5:0,6:0,7:0,8:0,9:0,10:0,11:0,12:0,13:0,14:0 switch (config) # ib baseqos sw0 vlarb-high i2=4:3 switch (config) # show ib baseqos sw0 vlarb-high 0:10,1:10,4:3,3:0,4:0,5:0,6:0,7:0,8:0,9:0,10:0,11:0,12:0,13:0,14:0</pre>	
Related Commands	show ib baseqos	
Notes	<ul style="list-style-type: none"> Unspecified elements will be filled with (index:0) You may have multiple entries with the same VL on this list 	

10.5.6.4.5 ib baseqos <port-type> vlarb-low <value>

	ib baseqos <port-type> vlarb-low {VW1 VW1 VW2 ...} no ib baseqos <port-type> vlarb-low Sets up to 15 VL to Weight mapping pairs for low priority processing. The no form of the command sets the attributes to their default settings.	
--	---	--

Syntax Description	port-type	<ul style="list-style-type: none"> ca—channel adapters rtr—routers sw0—ports 0 only of the switches swe—external ports of the switches
	VW[i]	<p>There are two possible options for this parameter:</p> <ul style="list-style-type: none"> A single vector (1 ...15) in the format of “#:#” separated by spaces, see example below. Format of “i#=X:Y” in order to change a specific entry (see example below)
Default	The default mapping is: 0:0,1:4,2:4,3:4,4:4,5:4,6:4,7:4,8:4,9:4,10:4,11:4,12:4,13:4,14:4	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # ib baseqos sw0 vlarb-low 1:1 switch (config) # show ib baseqos sw0 vlarb-low 1:1, 1:0, 2:0, 3:0, 4:0, 5:0, 6:0, 7:0, 8:0, 9:0, 10:0, 11:0, 12:0, 13:0, 14:0 switch (config) # ib baseqos sw0 vlarb-low i2=4:3 switch (config) # show ib baseqos sw0 vlarb-low 1:1, 1:0, 4:3, 3:0, 4:0, 5:0, 6:0, 7:0, 8:0, 9:0, 10:0, 11:0, 12:0, 13:0, 14:0</pre>	
Related Commands	show ib baseqos	
Notes	You may have multiple entries with the same VL on this list.	

10.5.6.4.6 ib baseqos reset-config

	ib baseqos reset-config Resets all basic QoS configuration options to defaults.	
Syntax Description	N/A	
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # ib baseqos reset-config</pre>	
Related Commands		
Notes		

10.5.6.4.7 show ib baseqos

	show ib baseqos <port-type> <baseqos-parameters> Displays the base IB QoS configuration.	
Syntax Description	port-type	<ul style="list-style-type: none"> ca—channel adapters rtr—routers sw0—ports 0 only of the switches swe—external ports of the switches

	baseqos-parameters	Possible values are: <ul style="list-style-type: none"> • high-limit—Display high limit (how many high pri before low) • max-vls—Display maximum number of VLs supported on CAs in subnet • sl2vl—Display current SL-to-VL mapping vector • vlarb-high—Display current high priority VL arbitration • vlarb-low—Display current low priority VL arbitration
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example	<pre>switch (config) # show ib baseqos ca high-limit 0</pre>	
Related Commands		
Notes		

10.5.6.4.8 ib qos

	<pre>ib qos no ib qos</pre> <p>Enables advanced QoS management on this node. The no form of the command disables advance QoS on this node.</p>
Syntax Description	N/A
Default	Advance QoS is disabled
Configuration Mode	config
History	3.1.0000
Example	<pre>switch (config) # show ib qos enable</pre>
Related Commands	show ib qos
Notes	

10.5.6.4.9 ib qos level

	<pre>ib qos level {<name> default} {mtu-limit <mtu> packet-life <time> pkey <number> rate-limit <rate-value> sl <sl-value> use <description>} no ib qos level {<name> default} {mtu-limit packet-life pkey rate-limit sl use}</pre> <p>Specifies a QoS level <name> or “default” parameters. The no form of the command set the parameters to default.</p>	
Syntax Description	<name> default	Specify a name for this qos group, or use the “default” for the default qos parameters
	mtu-limit <mtu>	MTU in bytes Possible values: 1k, 256, 2k, 4k, 512

	packet-life <time>	Time a packet can wait in switch egress queue before being dropped. The bytes from 4 microsecond up to 2 seconds or infinite. Possible values: 0-20 0–4usec 1–8usec ... 20–unlimited
	pkey <number>	PKEY value: ranges between -1 and 32767 (hex 0x7fff)
	rate-limit <rate-value>	Manages rate limits for QoS Policy levels Possible values (in Gbps): default, 2.5, 5, 10, 14, 20, 25, 40, 56, 100
	sl <sl-value>	Manages service level for QoS Policy levels Range: 0-15.
	use <description>	Specify usage description for this QoS level
Default		<ul style="list-style-type: none"> • use = “default QoS Level” • sl = 0 • mtu-limit = default • rate-limit = default • packet-life = 0x12 • pkey = -1
Configuration Mode	config	
History	3.1.0000	
	3.4.1100	Updated description of “rate-limit” parameter
Example	<pre>switch (config) # show ib qos my-qos-group my-qos-group: use = default QoS Level sl = 0 mtu-limit = 2K rate-limit = default packet-life = 0x12 pkey = -1</pre>	
Related Commands	show ib qos	
Notes		

10.5.6.4.10 ib qos match-rule

	ib qos match-rule <rule-index> {{destination source} <string> {pkey qos-class service-id} <index> {first last} <value>} qos-level-name <name> use <description>} no ib qos match-rule <rule-index> {{destination source} {pkey qos-class service-id} <index> {first last} } qos-level-name use} Manages QoS Policy match rules. The no form of the command set the QoS match-rule to default.	
Syntax Description	rule-index	Index of this match-rule Range: 0-4294967295
	destination source <string>	Manages destination or source for QoS Policy match rules
	pkey qos-class service-id <index>	Manages values for QoS Policy match rules

	{first last} <value>	First or last value range (per PKEY / qos-class of service ID)
	qos-level-name <name>	Name for the QoS level
	use <description>	Specify usage description for this QoS level
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # show ib qos match-rule 10 match-rule/10: match-rules: use = my-use match-rules: qos-level-name = DEFAULT</pre>	
Related Commands	show ib qos	
Notes		

10.5.6.4.11 ib qos port-group

	<p>ib qos port-group <name> {node-type <index> type <node-type> partition <name> pkey <number> port-guid <index> {first last} <value> port-name <index> name <name-value> use <description>}</p> <p>no ib qos port-group <name> {node-type <index> type partition pkey port-guid <index> {first last} port-name <index> name use }</p> <p>Manages QoS Policy port groups. The no form of the command removes a QoS port-group.</p>	
Syntax Description	<name>	Port group name
	node-type <index>	Node type index
	type <node-type>	A node type for this port group
	partition <name>	A Partition name
	pkey <number>	A PKEY number
	port-guid <index> {first last} <value>	Port-guid range
	port-name <index> name <name-value>	Port index name
	use <description>	Specify usage description for this QoS level
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config)# show ib qos port-group my-group port-group/my-group: port-groups: pkey = -1 port-groups: use = my-use</pre>	
Related Commands	show ib qos	
Notes		

10.5.6.4.12 ib qos ulp any

	ib qos ulp any {pkey service-id target-port-guid <index> {first last sl} <value> sl <sl-value>} no ib qos ulp any {pkey service-id target-port-guid <index> {first last sl} sl} Configures ULP any attributes. The no form of the command deletes ULP any attributes.	
Syntax Description	pkey <index>	Manages ULP default PKEY assignment
	service-id <index>	Manages default ULP Service ID match rule
	target-port-guid <index>	Manages ULP default target port GUID rule
	first last sl <value>	<ul style="list-style-type: none"> • first—first value in range • last—last value in range • sl—Service level for the ULP rule
	sl <sl-value>	Sets default SL
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib qos ulp any sl 2	
Related Commands	show ib qos	
Notes		

10.5.6.4.13 ib qos ulp ipoib

	ib qos ulp ipoib {default sl <sl-value> pkey <index> {first last sl} <value> } no ib qos ulp ipoib {default sl pkey <index>} Manages ULP IPOIB settings. The no form of the command deletes IPOIB settings.	
Syntax Description	default sl <sl-value>	Sets the default SL Range 1-15
	pkey <index>	Manages ULP default PKEY assignment
	first last sl <value>	<ul style="list-style-type: none"> • first—first value in range • last—last value in range • sl—service level for the ULP rule
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib qos ulp ipoib default sl 5	
Related Commands	show ib qos	
Notes		

10.5.6.4.14 `ib qos ulp <protocol-type>`

	<code>ib qos ulp <protocol-type> {default sl <sl-value> port-num< index> <first last sl> <value>}</code> <code>no ib qos ulp iser {default <sl> port-num1 <first last sl>}</code> Configures ULP IScsi Extensions for RDMA, Reliable Datagram Sockets or Sockets Direct Protocol attributes. The no form of the command deletes all rules.	
Syntax Description	protocol-type	iser—iSCSI extensions for RDMA (iSER) rds—reliable datagram sockets (RDS) sdp—sockets direct protocol (SDP)
	default sl <sl-value>	Sets the default SL Range 1-15
	port-num< index>	Port number index
	first last sl	<ul style="list-style-type: none"> • first—first in range • last—last in range • sl—service level for the ULP rule
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # ib qos ulp iser default sl 2</pre>	
Related Commands	show ib qos	
Notes		

10.5.6.4.15 `ib qos ulp srp`

	<code>ib qos ulp srp target-port-guid <index> <first last sl> <value></code> <code>no ib qos ulp srp target-port-guid <index></code> Configures Scsi Rdma Protocol attributes. The no form of the command deletes the rules.	
Syntax Description	target-port-guid <index>	The index of the target port GUID
	first last sl	<ul style="list-style-type: none"> • first—first in range • last—last in range • sl—service level for the ULP rule
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	<pre>switch (config) # ib qos ulp srp target-port-guid 1 sl 2</pre>	
Related Commands	show ib qos	
Notes		

10.5.6.4.16 show ib qos

	show ib qos [level match-rule port-group ulp] Displays InfiniBand QoS configurations	
Syntax Description	level	Displays QoS level configurations
	match-rule	Displays QoS match-rule configurations
	port-group	Displays QoS port-group configurations
	ulp	Displays QoS ulp configurations
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example	<pre>switch (config) # show ib qos level my-qos-level my-qos-level: use = my-use sl = 0 mtu-limit = 2K rate-limit = default packet-life = 0x12 pkey = -1</pre>	
Related Commands		
Notes		

10.5.6.5 Scatter Ports

10.5.6.5.1 ib sm scatter-ports

	ib sm scatter-ports <seed> no ib sm scatter-ports Activates scatter ports and sets seed for random number generation. The no form of the command deactivates the partition.	
Syntax Description	seed	Integer between 0-4294967295
Default	Disabled	
Configuration Mode	config	
History	3.6.8008	
Example	switch (config) # ib sm scatter-ports 123	
Related Commands	ib sm guid-routing-order-no-scatter	
Notes		

10.5.6.5.2 show ib sm scatter-ports

	show ib sm scatter-ports Displays scatter port seed.	
Syntax Description	N/A	
Default	N/A	

Configuration Mode	Any command mode
History	3.6.8008
Example	switch (config) # show ib sm scatter-ports Scatter ports seed: 234
Related Commands	ib sm scatter-ports
Notes	

10.5.6.6 GUID Routing Order

10.5.6.6.1 ib sm guid-routing-order add

	ib sm guid-routing-order add <guid> [position <pos>] Adds a new GUID to routing order list.	
Syntax Description	guid	GUID to add
	position	A position for the new GUID may be specified
Default	If no position is specified, the new GUID is added to the end of the list	
Configuration Mode	config	
History	3.6.8008	
Example	switch (config) # ib sm guid-routing-order add E4:1D:2D:03:00:3D:5E:87 position 6	
Related Commands	ib sm guid-routing-order-no-scatter	
Notes		

10.5.6.6.2 ib sm guid-routing-order delete

	ib sm guid-routing-order delete {<guid> position <pos>} Deletes a guid from routing order list. The guid can be chosen by its guid or by its position on guid routing order list.	
Syntax Description	guid	GUID to delete
	position	Deletes a GUID by specifying position number
Default	N/A	
Configuration Mode	config	
History	3.6.8008	
Example	switch (config) # ib sm guid-routing-order delete position 3 switch (config) # ib sm guid-routing-order delete E4:1D:2D:03:00:3D:5E:91	
Related Commands	ib sm guid-routing-order-no-scatter	
Notes		

10.5.6.6.3 `ib sm guid-routing-order move`

	<code>ib sm guid-routing-order move <guid> to-position <pos></code> Moves a GUID in the list to a specified position.	
Syntax Description	<code>guid</code>	GUID to move
	<code>position</code>	A position for the new GUID may be specified
Default	N/A	
Configuration Mode	config	
History	3.6.8008	
Example	<code>switch (config) # ib sm guid-routing-order move E4:1D:2D:03:00:3D:5E:91 to-position 2</code>	
Related Commands	<code>ib sm guid-routing-order-no-scatter</code>	
Notes		

10.5.6.6.4 `ib sm guid-routing-order move-down`

	<code>ib sm guid-routing-order move-down <guid></code> Moves a GUID position down in the GUID routing order list.	
Syntax Description	<code>guid</code>	GUID to move
Default	N/A	
Configuration Mode	config	
History	3.6.8008	
Example	<code>switch (config) # ib sm guid-routing-order move-down E4:1D:2D:03:00:3D:5E:91</code>	
Related Commands	<code>ib sm guid-routing-order-no-scatter</code>	
Notes		

10.5.6.6.5 `ib sm guid-routing-order move-up`

	<code>ib sm guid-routing-order move-up <guid></code> Moves a GUID position up in the GUID routing order list.	
Syntax Description	<code>guid</code>	GUID to move
Default	N/A	
Configuration Mode	config	
History	3.6.8008	
Example	<code>switch (config) # ib sm guid-routing-order move-up E4:1D:2D:03:00:3D:5E:91</code>	

Related Commands	ib sm guid-routing-order-no-scatter
Notes	

10.5.6.6.6 no ib sm guid-routing-order

	no ib sm guid-routing-order Disables the GUID routing order feature and cleans GUID routing order list.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.8008
Example	switch (config) # no ib sm guid-routing-order
Related Commands	ib sm guid-routing-order-no-scatter
Notes	

10.5.6.6.7 ib sm guid-routing-order-no-scatter

	ib sm guid-routing-order-no-scatter no ib sm guid-routing-order-no-scatter Enables randomization for destinations mentioned in GUID order list. The no form of the command disables randomization for destinations mentioned in GUID order list.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.8008
Example	switch (config) # ib sm guid-routing-order-no-scatter
Related Commands	ib sm guid-routing-order * ib sm scatter-ports
Notes	If scatter ports (randomization of the output port) is set to anything but zero, guid-routing-order-no-scatter defines whether or not randomization should be applied to the destination GUIDs mentioned in the GUID routing order list

10.5.6.6.8 show ib sm guid-routing-order

	show ib sm guid-routing-order Displays current GUID routing order list.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.8008

Example	<pre>switch (config) # show ib sm guid-routing-order 1: E4:1D:2D:03:00:3D:5E:85 2: E4:1D:2D:03:00:3D:5E:82 3: E4:1D:2D:03:00:3D:5E:81 4: E4:1D:2D:03:00:3D:5E:84 5: E4:1D:2D:03:00:3D:5E:86 6: E4:1D:2D:03:00:3D:5E:87 7: E4:1D:2D:03:00:3D:5E:90 8: E4:1D:2D:03:00:3D:5E:88 9: E4:1D:2D:03:00:3D:5E:83</pre>
Related Commands	ib sm guid-routing-order-no-scatter
Notes	

10.5.6.6.9 show ib sm guid-routing-order-no-scatter

	<pre>show ib sm guid-routing-order-no-scatter</pre> <p>Displays the status of the GUID-routing-order-no-scatter feature</p>
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.8008
Example	<pre>switch (config) # show ib sm guid-routing-order-no-scatter guid_routing_order_no_scatter: disabled</pre>
Related Commands	<pre>ib sm guid-routing-order * ib sm scatter-ports</pre>
Notes	

10.5.6.7 Bulk Update Mode

10.5.6.7.1 ib sm bulk-update enable

	<pre>ib sm bulk-update enable no ib sm bulk-update enable</pre> <p>Enables bulk update mode. The no form of the command disables bulk update mode.</p>
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.6.8008
Example	<pre>switch (config) # ib sm bulk-update enable</pre>
Related Commands	<pre>show ib partition show ib sm bulk-update</pre>
Notes	

10.5.6.7.2 show ib sm bulk-update

	show ib sm bulk-update Displays the status of bulk-update mode.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.6.8008
Example	<pre>switch (config) # show ib sm bulk-update ib sm bulk-update: enabled</pre>
Related Commands	ib sm bulk-update enable
Notes	

10.5.6.8 ibdiagnet

10.5.6.8.1 ibdiagnet

	ibdiagnet [parameters]	
Syntax Description	parameters	ibdiagnet native parameters
Default	N/A	
Configuration Mode	Any command mode	
History	3.9.3100	
Example	<pre>switch (config) # ibdiagnet</pre>	
Related Commands	show ibdiagnet file ibdiagnet upload file ibdiagnet delete	
Notes	To know the optional parameters, run ibdiagnet -h.	

10.5.6.8.2 show ibdiagnet

	show ibdiagnet Show output from latest call to ibdiagnet
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.9.3100
Example	<pre>switch (config) # show ibdiagnet</pre>
Related Commands	ibdiagnet
Notes	

10.5.6.8.3 file ibdiagnet upload

	file ibdiagnet upload <file name> <upload_url> Upload ibdiagnet archive of output files (from latest call to ibdiagnet) to a remote host.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.9.3100
Example	<pre>switch (config) # file ibdiagnet upload ibdiagnet_output.gz scp:// username:password@192.168.10.125/var/www/html/<image_name></pre>
Related Commands	ibdiagnet file ibdiagnet delete
Notes	

10.5.6.8.4 file ibdiagnet delete


	file ibdiagnet delete <file name> Deletes the specified ibdiagnet archive file.
Syntax Description	<file name> File name
Default	N/A
Configuration Mode	Any command mode
History	3.9.3100
Example	<pre>switch (config) # file ibdiagnet delete ibdiagnet_output.gz</pre>
Related Commands	file ibdiagnet upload
Notes	

10.6 Subnet Manager High Availability



All nodes in an SM HA subnet must be of the same CPU type (e.g. x86), and must run the same MLNX-OS version.

High availability (HA) refers to a system or component that is continuously operational for a desirably extended period of time.



NVIDIA Subnet Manager (SM) HA reduces subnet downtime and disruption as it is continuously operational for a desirably long length of time. It assures continuity of the work even when one of the SMs dies. The database is synchronized with all the nodes participating in the InfiniBand subnet and a configuration change is prepared. The synchronization is done out-of-band using an Ethernet management network.

NVIDIA SM HA allows the systems' manager to enter and modify all InfiniBand SM configuration of different subnet managers from a single location. It creates an InfiniBand subnet and associates all the NVIDIA management appliances that are attached to the same InfiniBand subnet into that InfiniBand subnet ID. All subnet managers can be controlled, started, or stopped from this address.

All the nodes that participate in the NVIDIA SM HA are joined to the InfiniBand subnet ID and once joined, the synchronized SMs are launched. One of the nodes is elected as Master and the others are Slaves (or down). NVIDIA SM HA uses an IP address (VIP) that is always directed to the SM HA master to monitor the SM state and to verify that all configurations are executed.

10.6.1 Joining, Creating or Leaving an InfiniBand Subnet ID

When transitioning from standalone into a group or vice versa, a few seconds are required for the node state to stabilize. During that time, group feature commands (e.g. SM HA commands) should not be executed. To run group features, wait for the CLI prompt to turn into [standalone:master], [<group>:master] or [<group>:standby] instead of [standalone:*unknown*] or [<group>:*unknown*].

An InfiniBand subnet is formed by a network of InfiniBand nodes interconnected via InfiniBand switches. It includes all systems that can run an SM and is part of the SM HA domain. A switch that can potentially run an SM must be a member of an InfiniBand subnet ID to be associated with the NVIDIA SM HA domain. An IB subnet is recognized by its ID which is used by the system to either join or leave the subnet.

Every system that is not associated to an existing IB subnet (has never been part of an IB subnet or has left an existing one) or does not have MLNX-OS license installed, is by default associated to a subnet called "Standalone".

In order to create, join or leave an InfiniBand subnet, one may use the following commands:

- Create - "ib ha <IB_subnet_ID> ip <ip_addr> <netmask>"

- Join - “ib ha <IB_subnet_ID>”
- Leave - “no ib ha”

When leaving an SM HA cluster, SM configuration is not saved on the node leaving the cluster. After leaving, the configuration is reset to its default values.

For further information see section [“Creating and Adding Systems to an InfiniBand Subnet ID”](#).

10.6.2 MLNX-OS Management Centralized Location

MLNX-OS centralized management infrastructure enables the user to configure or modify an existing configuration and monitor the subnet running status. MLNX-OS centralized management IP (VIP) is defined when a new subnet manager is created by running the command “ib ha <IB_subnet_ID> ip <ip_addr> <netmask>”. The created VIP is used as the current subnet master’s alias thus, assumes the same roles as the master.

The VIP always points to one of the systems part of the SM HA domain. It is always active even if one or more of the members are down. For example:

```
switch (config) # ib ha subnet2 ip 192.168.10.110 255.255.255.0
```

10.6.3 High Availability Node Roles

A node is an InfiniBand switch system. Every node member of an IB subnet ID has one of the following roles:

- Master - the node that manages SM configurations and provides services to the Virtual IP (VIP) addresses
- Standby - the node that replaces the Master node and takes over its responsibilities once the Master node is down
- Offline - has run an SM in the past and is currently offline, or it was created manually by the “ib smnode <node name> create” command. If the node has been removed from the environment, you can remove it from the list with the “no ib smnode xxx” command.

To see the mode of the current node, look at the CLI prompt for the following format:

```
<host name> [<subnet ID>:<mode>] [standalone: master] (config) #
```

For example:

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

To see a list of the existing nodes and details about the running state, run the command “show ib smnodes {brief}”.

10.6.4 Configuring MLNX-OS SM HA Centralized Location

The IP is used to configure or modify the existing configuration and monitor the subnet running status. To configure your IP, run the command “ib ha <IB_subnet_ID> ip <ip_addr> <netmask>”:

```
switch [standalone: master] (config) # ib ha subnet2 ip 192.168.10.110 255.255.255.0
switch [subnet2: master] (config) #
```

10.6.5 Creating and Adding Systems to an InfiniBand Subnet ID

To create and add systems to a subnet:

1. Log into the system from which you intend to create the subnet.
2. Enter config mode. Run:

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

3. Create a new subnet using the command “ib ha <IB_subnet_ID> ip <ip_addr> <netmask>”.
Run:

```
switch [standalone: master] (config) # ib ha subnet2 ip 192.168.10.110 255.255.255.0
switch [subnet2: master] (config) #
```

You must run the “ib ha <IB_subnet_ID> ip <ip_addr> <netmask>” command only once per subnet ID.

4. Log into the system that you are going to join to the new created subnet.
5. Join the system to the subnet, using the “ib ha <IB_subnet_ID>” command. Run:

```
switch [standalone: master] (config) # ib ha subnet2
switch [subnet2: standby] (config) #
```

10.6.6 Restoring Subnet Manager Configuration

In instances where the SM configuration becomes corrupted or the subnet manager cannot raise any logical links it is suggested that you restore the default SM configuration.

To restore subnet manager configuration:

1. Enter config mode. Run:

```
*switch [subnet2: master] > enable
*switch [subnet2: master] # configure terminal
*switch [subnet2: master] (config) #
```

2. Run the command “ib sm reset-config”. Run:

```
*switch [subnet2: master] (config) # ib sm reset-config
```

The asterisk in the example above (*switch) indicates the local system from where the command is running.

In order to receive information on the running state of a specific node one could run one of the following commands with its requested parameter:

- show ib smnode <name> sm-running
- show ib smnode <name> sm-state
- show ib smnode <name> sm-priority
- show ib smnode <name> active
- show ib smnode <name> ha-state
- show ib smnode <name> ha-role

10.6.6.1 Subnet Manager Configuration

To configure the subnet manager, log into the centralized management IP (VIP). Once the SM configuration is created, the SM database is duplicated to the other nodes.

The SM must be configured from MLNX-OS centralized management IP (VIP). All the configurations that are not created or modified in the master node (using the VIP) are overridden by the master configuration.

The user can configure different SM parameters such as where to run the SM(s) or the SM priority by running the commands according to the desired action.

10.6.6.2 NVIDIA High Availability and OpenSM Handover/Failover

NVIDIA products are fully compliant and interoperable with OpenSM.

Once an SM fails, the SM which takes over the subnet needs to reproduce the internal state of the failed master. Most of the information required is obtained by scanning the subnet and extracting the information from the devices. However, some information which is not stored directly in the network devices cannot be reproduced this way. InfiniBand management architecture limits such information to data exchanged between clients (either user-level programs or kernel modules) and the Subnet Administration (SA) service (attached to the SM). The SA keeps this set of client registrations in an internal data structure called SA-DB. The SA-DB information includes the multicast groups, the multicast group members, subscriptions for event forwarding and service records.

The new SM may retrieve the SA-DB by requesting the clients to re-register with the SA or by obtaining a copy of the previous master SM internal SA-DB via an SA-DB dump file. The client-re-registration offers database correctness and the SA-DB dump file replication provides lower setup time. Client re-registration is required since the SA-DB may not be up-to-date on the registrations listed in the master SM.

Furthermore, since the SM does not maintain SA-DB information for unknown nodes, it is very possible that some of the SA-DB information relating to nodes momentarily disconnected from the master SM become purged. Therefore, these nodes must re-register with the new SM when they are

reconnected (they receive a client-re-register request from the SM). Relying only on client re-registration is also non-optimal as it takes some time to recreate the entire SA-DB and the network state.

NVIDIA SM HA replicates the SA-DB dump file from the current master SM to all the standby SMs running on NVIDIA switches. The SA-DB dump file replication provides further optimization to the standby SM that becomes master.

Standby SM loads the existing SA-DB file the old master has used. By using the existing SA-DB the amount of processing needed on client re-registration is lessened resulting in a reduced time to complete setting up the network.

SM HA does not replace InfiniBand spec requirement for client re-registration.

When running an SM HA cluster with more than 2 active OpenSM instances, IB multicast applications need to support client re-register or they may not work correctly after OpenSM failover.

10.6.7 SM HA Commands

10.6.7.1 ib ha

	<code>ib ha <IB_subnet_ID> [ip <IP address> <subnet mask> [force]]</code> <code>no ib ha</code> Creates a subnet <IB_subnet_ID> with the specified IP. The no form of the command removes this node from an InfiniBand subnet ID.	
Syntax Description	IB subnet ID	Simple group name for shared IB config
	ip <IP address>	Assigns management IP address
	netmask	Netmask (e.g. 255.255.255.0 or /24)
	force	Joins if exists or creates if not
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	<code>switch (config) # ib ha my-subnet</code>	
Related Commands	show ib ha	
Notes	A new subnet may be joined only after leaving the current one	

10.6.7.2 ib smnode

	<code>ib smnode <hostname> [create disable enable sm-priority <priority>]</code> <code>no ib smnode <hostname> [create disable enable sm-priority]</code> Manages HA SM. The no form of the command removes HA SM node configuration.	
Syntax Description	hostname	Specifies <hostname> SM configuration to modify.

	create	Creates SM configuration for selected node.
	disable	Makes SM inactive on selected node.
	enable	Makes SM active on selected node.
	sm-priority <priority>	Sets SM selected node priority (0=low, 15=high).
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config) # ib smnode switch-1133ce create	
Related Commands	show ib smnode show ib smnodes	
Notes		

10.6.7.3 show ib smnode

	show ib smnode <hostname> {active ha-role ha-state ip sm-priority sm-running sm-state} Displays SM High availability information.	
Syntax Description	hostname	Specifies <hostname> SM configuration to display
	active	Displays whether <hostname> is currently active
	ha-role	Displays the High Availability role of <hostname>. Possible return values are: offline, unknown, master, standby, or disabled
	ha-state	Possible return values are: offline, init, searching, joining, online, creating, waiting, leaving, join-sync, failed, removed, or regroup
	ip	Displays the local management IP address associated with the active node, <hostname>. If <hostname> is not active, the command displays "offline"
	sm-priority	Displays the SM priority for SM running on <hostname>
	sm-running	Displays if <hostname> has an SM running. The command will display "active" (that is, SM is running) only if <hostname> is currently active, has a license, is enabled as a potential SM, is active as SM, and if there is a maximum of 2 SMs in the fabric.
	sm-state	Displays if SM is enabled to run on <hostname>
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.8.1000	Updated Syntax Description
Example	switch (config) # show ib smnode my-hostname sm-state enabled	
Related Commands	show ib smnodes	
Notes		

10.6.7.4 show ib smnodes

	show ib smnodes [brief] Displays SM High availability information.	
Syntax Description	brief	Displays information on all HA nodes
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.8.1000	Updated example
	3.9.3100	Updated output to reflect the OpenSM master also when the command is triggered from non-SM master
Example	<pre>switch (config) # show ib smnodes HA state of switch infiniband-default: IB Subnet HA name: Mantaray142 HA IP address : 10.7.145.141/24 Active HA nodes : 2 HA node local information: Name : Mantaray142 (active) <--- (local node) SM-HA state: standby SM Running : stopped SM Enabled : enabled - master SM Priority: 0 IP : 10.7.144.142 HA node local information: Name : Mantaray141 (active) SM-HA state: master SM Running : stopped SM Enabled : disabled SM Priority: 0 IP : 10.7.144.141</pre>	
Related Commands		
Notes		

show ib ha

	show ib ha [brief] Displays information about all the systems that are active or might be able to run SM.	
Syntax Description	brief	Displays brief HA information
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
	3.9.1000	Updated example

Example	<pre> switch (config) # show ib ha Global HA state: IB Subnet HA name: Barracuda-s HA IP address : 10.7.48.100/24 Active HA nodes : 2 HA node local information: Name : barracuda-216 (active) <--- (local node) SM-HA state : standby IP : 10.7.48.50 Virtual switch membership: infiniband-default HA node local information: Name : barracuda-217 (not active) IP : offline Virtual switch membership: infiniband-default HA node local information: Name : scorpionib2-19 (active) SM-HA state : master IP : 10.7.51.169 Virtual switch membership: infiniband-default switch (config) # show ib ha brief Global HA state: IB Subnet HA name: Barracuda-s HA IP address : 10.7.48.100/24 Active HA nodes : 3 ----- ID Local node SM-HA state IP Virtual switch membership ----- barracuda-216 * standby 10.7.48.50 infiniband-default barracuda-217 standby 10.7.48.51 infiniband-default scorpionib2-19 master 10.7.51.169 infiniband-default </pre>
Related Commands	
Notes	

11 Appendixes

The document contains the following appendixes:

- [Appendix: Enhancing System Security According to NIST SP 800-131A](#)
- [Appendix: Splunk Integration with NVIDIA Products](#)
- [Appendix: Show Commands Not Supported By JSON API](#)

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



Our switch systems, by default, work with NIST SP 800-131A, as described in the table below.

This appendix describes how to enhance the security of a system in order to comply with the NIST SP 800-131A standard. This standard is a document which defines cryptographically “acceptable” technologies. This document explains how to protect against possible cryptographic vulnerabilities in the system by using secure methods. Because of compatibility issues, this security state is not the default of the system and it should be manually set.

Some protocols, however, cannot be operated in a manner that complies with the NIST SP 800-131A standard.

Component	Configuration	Command
HTTP	HTTP disabled	no web http enable
HTTPS	HTTPS enabled	no web https enable
	SSL ciphers = TLS1.2	web https ssl ciphers all
	SSL renegotiation disabled	web https ssl renegotiation enable
SSH	SSH version = 2	ssh server min-version 1
	SSH ciphers = aes256-ctr, aes192-ctr, aes128-ctr, aes128-gcm@openssh.com, aes256-gcm@openssh.com	no ssh server security strict

11.1.1 Web Certificate

The OS supports signature generation of sha256WithRSAEncryption, sha1WithRSAEncryption self-signed certificates, and importing certificates as text in PEM format.

To configure a default certificate:

1. Create a new sha256 certificate.

```
switch (config) # crypto certificate name <cert name> generate self-signed hash-algorithm sha256
```

For more details and parameters refer to the command [“crypto certificate name”](#).

2. Show crypto certificate detail.

```
switch (config) # show crypto certificate detail
```

Search for “signature algorithm” in the output.

3. Set this certificate as the default certificate. Run:

```
switch (config) # crypto certificate default-cert name <cert name>
```

To configure default parameters and create a new certificate:

1. Define the default hash algorithm.

```
switch (config) # crypto certificate generation default hash-algorithm sha256
```

2. Generate a new certificate with default values.

```
switch (config) # crypto certificate name <cert name> generate self-signed
```

When no options are selected, the generated certificate uses the default values for each field.

To test strict mode connect to the WebUI using HTTPS and get the certificate. Search for “signature algorithm”.

There are other ways to configure the certificate to sha256. For example, it is possible to use “certificate generation default hash-algorithm” and then regenerate the certificate using these default values.

It is recommended to delete browsing data and previous certificates before retrying to connect to the WebUI.

Make sure not to confuse “signature algorithm” with “Thumbprint algorithm”.

11.1.2 SNMP

SNMPv3 supports configuring username, authentication keys and privacy keys. For authentication keys it is possible to use MD5 or SHA. For privacy keys AES or DES are to be used.

To configure strict mode, create a new user with HMAC-SHA1-96 and AES-128. Run:

```
switch (config) # snmp-server user <username> v3 auth sha <password1> priv aes-128 <password2>
```

To verify the user in the CLI, run:

```
switch (config) # show snmp user
```

To test strict mode, configure users and check them using the CLI, then run an SNMP request with the new users.

SNMPv1 and SNMPv2 are not considered to be secure. To run in strict mode, only use SNMPv3.

11.1.3 HTTPS

By default, the OS supports HTTPS encryption using TLS1.2 only. Working in TLS1.2 mode also bans MD5 ciphers which are not allowed per NIST 800-131a. In strict mode, the switch supports encryption with TLS1.2 only with the following supported ciphers:

- RSA_WITH_AES_128_CBC_SHA256
- RSA_WITH_AES_256_CBC_SHA256
- DHE_RSA_WITH_AES_128_CBC_SHA256
- DHE_RSA_WITH_AES_256_CBC_SHA256
- TLS_RSA_WITH_AES_128_GCM_SHA256
- TLS_RSA_WITH_AES_256_GCM_SHA384
- TLS_DHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_DHE_RSA_WITH_AES_256_GCM_SHA384

To enable all encryption methods, run:

```
switch (config) # web https ssl ciphers all
```

To enable only TLS ciphers (enabled by default), run:

```
switch (config) # web https ssl ciphers TLS
```

To enable HTTPS strict mode, run:

```
switch (config) # web https ssl ciphers TLS1.2
```

To verify which encryption methods are used, run:

```
switch (config)# show web
Web User Interface:
Web interface enabled: yes
HTTP enabled: yes
HTTP port: 80
HTTP redirect to HTTPS: no
HTTPS enabled: yes
HTTPS port: 443
HTTPS ssl-ciphers: TLS1.2
HTTPS certificate name: default-cert
Listen enabled: yes
No Listen Interfaces.
```

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

Web file transfer proxy:
Proxy enabled: no

Web file transfer certificate authority:
HTTPS server cert verify: yes
HTTPS supplemental CA list: default-ca-list
```

On top of enabling HTTPS, to prevent security breaches HTTP must be disabled.

To disable HTTP, run:

```
switch (config) # no web http enable
```

11.1.4 Code Signing

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

Strict mode is operational by default.

11.1.5 SSH

The SSH server on the switch by default uses secure ciphers only, message authentication code (MAC), key exchange methods, and public key algorithm. When configuring SSH server to strict mode, the aforementioned security methods only use approved algorithms as detailed in the NIST 800-181A specification and the user can connect to the switch via SSH in strict mode only.

To enable strict security mode, run the following:

```
switch (config) # ssh server security strict
```

The following ciphers are disabled for SSH when strict security is enabled:

- 3des-cbc
- aes256-cbc
- aes192-cbc
- aes128-cbc
- rijndael-cbc@lysator.liu.se

The no form of the command disables strict security mode.

Make sure to configure the SSH server to work with minimum version 2 since 1 is vulnerable to security breaches.

To configure min-version to strict mode, run:

```
switch (config) # ssh server min-version 2
```

Once this is done, the user cannot revert back to minimum version 1.

11.1.6 LDAP

By default, the switches support LDAP encryption SSL version 3 or TLS1.0 up to TLS1.2. The only banned algorithm is MD5 which is not allowed per NIST 800-131a. In strict mode, the switch supports encryption with TLS1.2 only with the following supported ciphers:

- DHE-DSS-AES128-SHA256
- DHE-RSA-AES128-SHA256
- DHE-DSS-AES128-GCM-SHA256
- DHE-RSA-AES128-GCM-SHA256
- DHE-DSS-AES256-SHA256
- DHE-RSA-AES256-SHA256
- DHE-DSS-AES256-GCM-SHA384
- DHE-RSA-AES256-GCM-SHA384
- ECDH-ECDSA-AES128-SHA256
- ECDH-RSA-AES128-SHA256
- ECDH-ECDSA-AES128-GCM-SHA256
- ECDH-RSA-AES128-GCM-SHA256
- ECDH-ECDSA-AES256-SHA384
- ECDH-RSA-AES256-SHA384
- ECDH-ECDSA-AES256-GCM-SHA384
- ECDH-RSA-AES256-GCM-SHA384
- ECDHE-ECDSA-AES128-SHA256
- ECDHE-RSA-AES128-SHA256
- ECDHE-ECDSA-AES128-GCM-SHA256
- ECDHE-RSA-AES128-GCM-SHA256
- ECDHE-ECDSA-AES256-SHA384
- ECDHE-RSA-AES256-SHA384
- ECDHE-ECDSA-AES256-GCM-SHA384
- ECDHE-RSA-AES256-GCM-SHA384
- AES128-SHA256
- AES128-GCM-SHA256
- AES256-SHA256
- AES256-GCM-SHA384

To enable LDAP strict mode, run the following:

```
switch (config) # ldap ssl mode {start-tls | ssl}
```


Both modes operate using SSL. The difference lies in the connection initialization and the port used.

11.2 Appendix: Splunk Integration with NVIDIA Products



Splunk automatically clusters millions of log records in real time back into their patterns and finds connections between those patterns to form the baseline flows of each software individually, thus enabling you to search, monitor and analyze that data to discover powerful insights across multiple use cases.

This appendix provides a guide on the first steps with Splunk and helps you to begin enjoying reduced time in detecting and resolving production problems.

11.2.1 Getting Started with Splunk

1. Download Splunk and extract the Splunk Enterprise version. (Splunk software is available as an RPM or TGZ.)

2. Create a Splunk User /group. Run:

```
[root@server] groupadd splunk
[root@server] useradd -d /opt/splunk -m -g splunk splunk
```

3. Splunk installation. Run:

```
[root@server] tar -xzf splunk-7.0.0-c8a78efdd40f-Linux-x86_64.tgz
[root@server] ls
```

4. A new folder called Splunk is created.

```
[root@server] cp -rp splunk/* /opt/splunk/
[root@server] chown -R splunk: /opt/splunk/
[root@server] su - splunk
[splunk@server] cd bin
[splunk@server] ./splunk start --accept-license
```

Now you can access your Splunk WebUI at <http://IP:8000/> or <http://hostname:8000/>. You need to make sure that port 8000 is open in your server firewall.

11.2.2 Switch Configuration

In this example we are not using the default UDP port 514 to show that any other port can be also used.

5. In order to add a task, the switch must be configured to send logs to our Splunk server. Run:

```
switch > enable
switch # configure terminal
switch (config) # show snmp
SNMP enabled:      yes
SNMP port:         161
System contact:
```

```

System location:

Read-only communities:
  public

Read-write communities:
  (none)

Interface listen enabled: yes
No Listen Interfaces.

switch (config) # snmp-server host 10.212.23.1 informs port 8597
switch (config) # snmp-server host 10.212.23.1 traps port 8597
switch (config) # snmp host 10.212.23.1 informs 8597
switch (config) # snmp host 10.212.23.1 traps 8597


Summary configuration:

switch (config) # show running-config
## Logging configuration
##
logging 10.212.23.1
logging 10.212.23.1 port 8597
logging 10.212.23.1 trap info
logging 10.212.23.1 trap override class events priority err
logging monitor events notice
logging receive
## SNMP configuration
no snmp-server host 10.209.21.221 disable
snmp-server host 10.209.21.221 traps port 8597 version 2c
no snmp-server host 10.212.23.1 disable
snmp-server host 10.212.23.1 traps port 8597 version 2c 8597

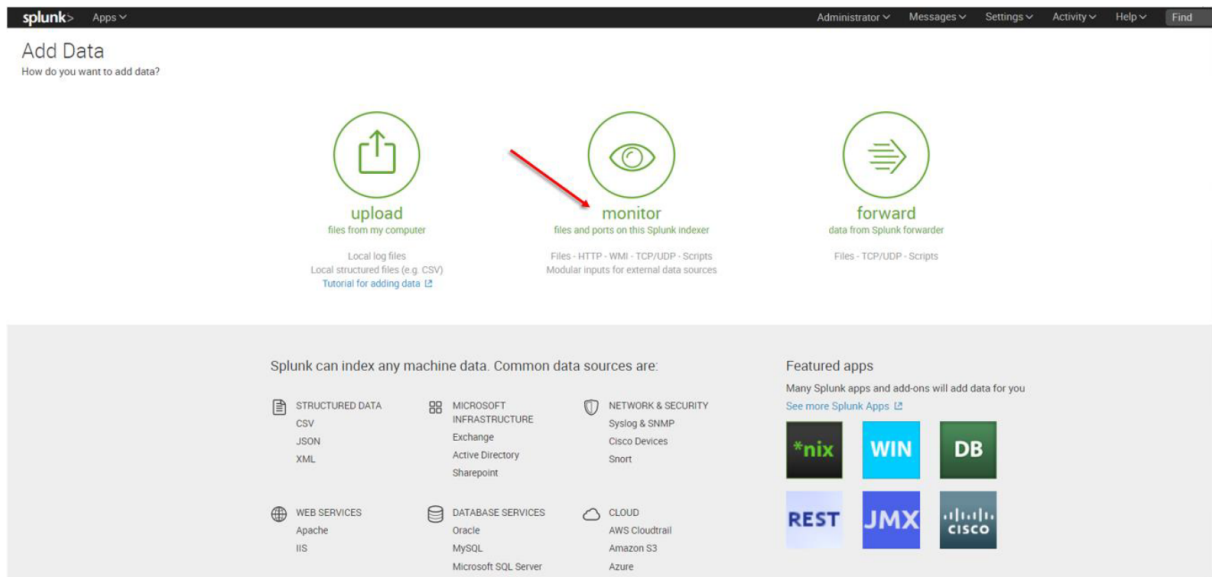
```

11.2.3 Adding a Task

6. The first screen encountered after signing into the Splunk WebUI includes the “Add Data” icon.



7. The “Add Data” tab opens up with three options: Upload, Monitor, and Forward. Here our task is to monitor a folder, so we click Monitor. to proceed

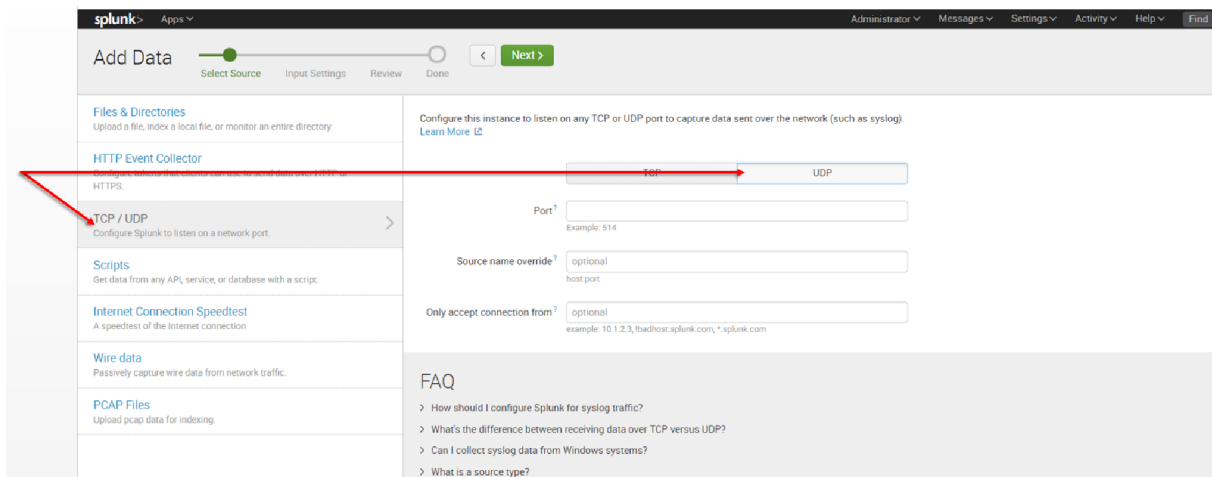


In the Monitor option, the following four categories are available:

- File & Directories - monitor files/folders
- HTTP Event Collector - monitor data streams over HTTP
- TCP/UDP - monitor service ports
- Scripts - monitor scripts


11.2.4 Retrieving Data from TCP and UDP Ports

8. Per our current purpose, we choose TCP/UDP option.

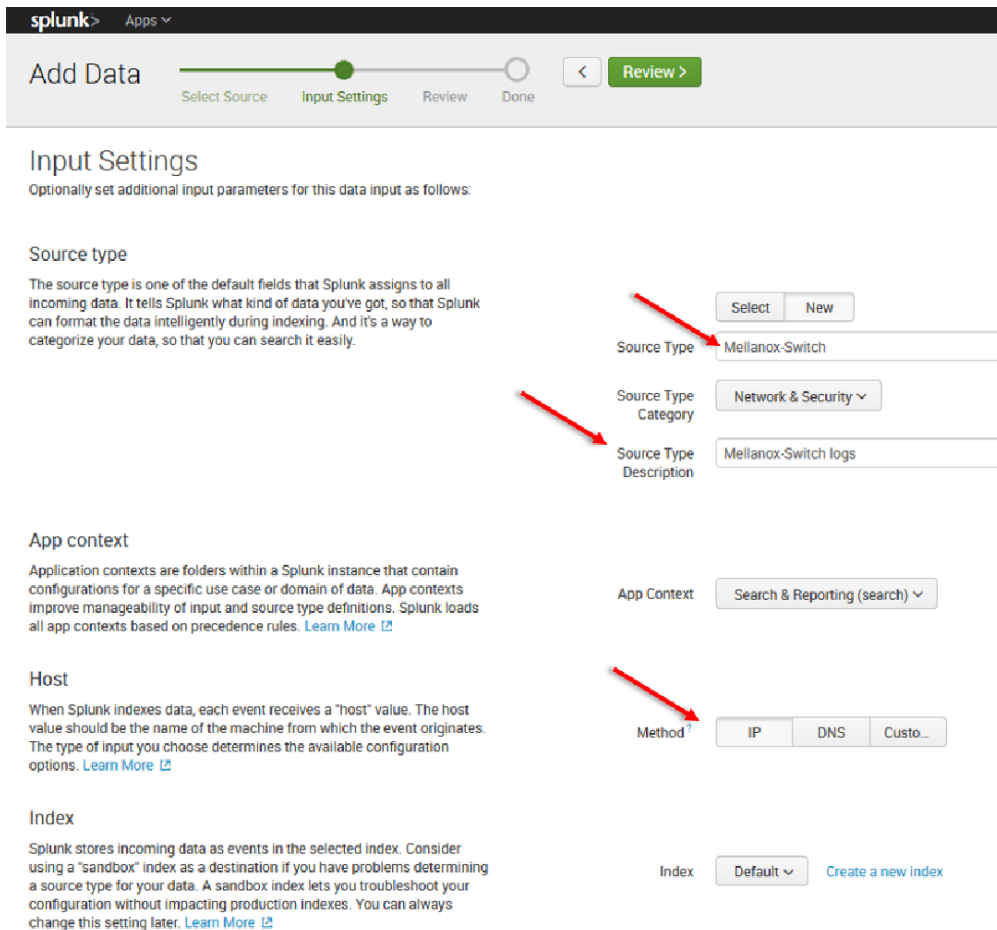


9. Click the TCP or UDP button to choose between a TCP or UDP input, and enter a port number in the “Port” field.

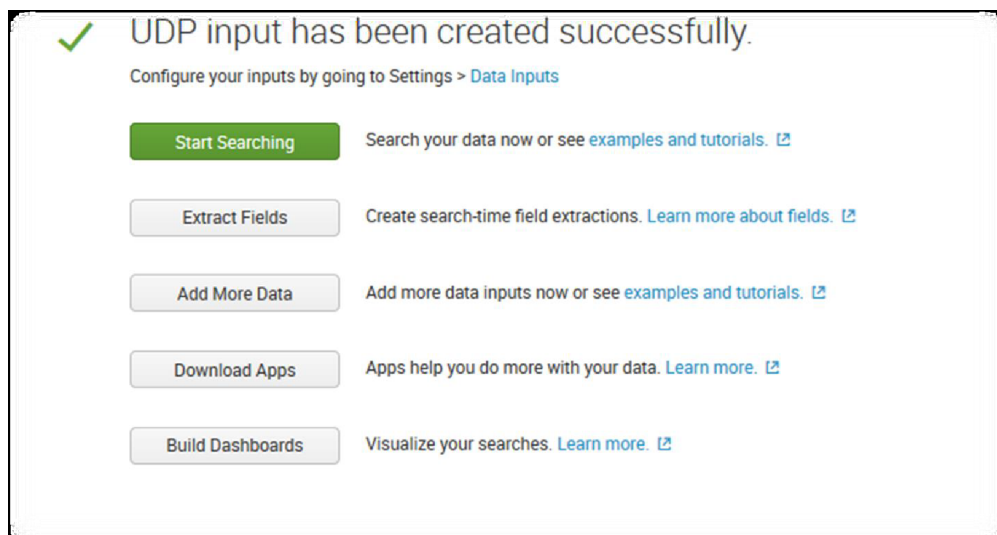
10. In the “Source name override” field, enter a new source name to override the default source value, if required.



11. Click “Next” to continue to the Input Settings page where we will create a new source type called Mellanox-Switch.



12. Click Next > Review > Done > Start Searching



11.2.5 SNMP Input to Poll Attribute Values and Catch Traps

SNMP represents an incredibly rich source of data that you can get into Splunk for visibility across a very diverse IT landscape.

SNMP agents may also send notifications, called Traps, to an SNMP trap listening daemon.

11.2.5.1 Getting Started

Browse to Splunkbase and download the SNMP Modular Input from <https://splunkbase.splunk.com/app/1537/>.

To install, simply untar the file to `SPLUNK_HOME/etc/apps` and restart Splunk.

11.2.5.2 Configuration

Login to the Splunk WebUI and go to Manager > Add Data > Monitor > SNMP > New, and set up your input data.

splunk> Apps

Add Data

Select Source Input Settings Review Done

< Next >

- Files & Directories**
Upload a file, index a local file, or monitor an entire directory.
- HTTP Event Collector**
Configure tokens that clients can use to send data over HTTP or HTTPS.
- TCP / UDP**
Configure Splunk to listen on a network port.
- Scripts**
Get data from any API, service, or database with a script.
- SNMP**
SNMP input to poll attribute values and catch traps
- Internet Connection Speedtest**
A speedtest of the Internet connection
- Wire data**
Passively capture wire data from network traffic.
- PCAP Files**
Upload pcap data for indexing.



splunk> Apps Administrator Messages Settings Activity Help

Add Data

Select Source Done < Next >

- Files & Directories
- HTTP Event Collector
- TCP / UDP
- Scripts
- SNMP** >
- Internet Connection Speedtest
- Wire data
- PCAP Files

Response Handler arguments string, key=value,key2=value2

SNMP Attribute polling settings

Destination:
IP or hostname of the device you would like to query, or a comma delimited list

Port:
The SNMP port. Defaults to 161

Object Names List:
1 or more Objects Names, comma delimited, in either textual(iso.org.dod.internet.mgmt.mib-2.system.sysDescr.0) or numerical(1.3.6.1.2.1.1.3.0) format

Interval:
How often to run the SNMP query (in seconds). Defaults to 60 seconds

Perform GET BULK:
Whether or not to perform a SNMP GET BULK operation. This will retrieve all the object attributes in the sub tree of the declared OIDs. Be aware of potential performance issues. http://www.net-snmp.org/wiki/index.php/GETBULK. Defaults to false.

Perform GET SUBTREE:
Whether or not to perform a SNMP GET SUBTREE operation. This will retrieve all the object attributes in the sub tree of the declared OIDs. Be aware of potential performance issues. http://www.net-snmp.org/wiki/index.php/GETNEXT. Defaults to false.

Split Bulk Results:
Whether or not to split up bulk output into individual events. Defaults to false.

Non Repeaters (for GET BULK):
The number of objects that are only expected to return a single GETNEXT instance, not multiple instances. Managers frequently request the value of sysUpTime and only want that instance plus a list of other objects. Defaults to 0.

Max Repetitions (for GET BULK):
The number of objects that should be returned for all the repeating OIDs. Agent's must truncate the list to something shorter if it won't fit within the max-message size supported by the command generator or the agent. Defaults to 25.

Source type

Source type
Set sourcetype field for all events from this source.

Set sourcetype

Select source type from list

Splunk classifies all common data types automatically, but if you're looking for something specific, you can find more source types in the Splunkbase apps browser or online at www.splunkbase.com.

More settings

Host
Host field value

Index
Set the destination index for this source.
Index

13. After configuration is complete it is recommend to run MellanoX-Switch again: Search > Data Summary > Sourcetypes > MellanoX-Switch.

The screenshot shows the Splunk Search & Reporting interface. A 'Data Summary' window is open, displaying a table of source types. The 'Sourcetypes (2)' tab is selected, and 'MellanoX-Switch' is highlighted in the table. A red arrow points from the 'MellanoX-Switch' entry in the table to the 'Add to search' button in the search bar of the next screenshot.

Sourcetype	Count	Last Update
MellanoX-Switch	1,278,154	10/18/17 8:53:28.000 AM

14. Select "MellanoX-Switch" and "Add to search".

The screenshot shows the Splunk Search & Reporting interface with a search query: `ERR OR [sourcetype=access_* (404 OR 500 OR 503)]`. The search results are displayed in a timeline view, showing 540 events. A red arrow points from the 'Add to search' button in the search bar to the 'Add to search' button in the search results table.

i	Time	Event
> 1	10/18/17 9:32:13.000 AM	Oct 18 09:32:13 10.209.21.221 had[5058]: [had.INFO]: had_handle_query_request: err=0 index = main linecount = 1 source = 10.209.21.221:8597 sourcetype=MellanoX-Switch
> 2	10/18/17 9:32:13.000 AM	Oct 18 09:32:13 10.209.21.221 had[5058]: [had.INFO]: had_handle_query_request: err=0 index = main linecount = 1 source = 10.209.21.221:8597 sourcetype=MellanoX-Switch
> 3	10/18/17 9:32:13.000 AM	Oct 18 09:32:13 10.209.21.221 chad[4659]: [chad.INFO]: chad_handle_query_request: err=0 index = main linecount = 1 source = 10.209.21.221:8597 sourcetype=MellanoX-Switch
> 4	10/18/17 9:31:55.000 AM	Oct 18 09:31:55 10.209.21.221 health[4753]: [health.INFO]: health_check_perform_sweep, err=0 index = main linecount = 1 source = 10.209.21.221:8597 sourcetype=MellanoX-Switch

15. You can add to search any value that is relevant for you.



nat 20 Per Page v

Event	
Oct 18 09:01:31 10.209.21.221	dhclient[4508]: dhc6: send_packet6() sent -1 of 151 bytes <small>host = 10.209.21.221 linecount = 1 source = 10.209.21.221:8597 sourcetype = Mellanox-Switch</small>
Oct 18 09:01:31 10.209.21.221	dhclient[4508]: send_packet6: Network is unreachable <small>host = 10.209.21.221 linecount = 1 source = 10.209.21.221:8597 sourcetype = Mellanox-Switch</small>
Oct 18 09:01:31 10.209.21.221	dhclient[4508]: XMT: Solicit on mgmt1, interval 109220ms. <small>host = 10.209.21.221 linecount = 1 source = 10.209.21.221:8597 sourcetype = Mellanox-Switch</small>
Oct 18 09:01:31 10.209.21.221	arpd[4965]: TID 140429637707520: [arpd.INFO]: linux_ifindex: 4 <small>host = 10.209.21.221 linecount = 1 source = 10.209.21.221:8597 sourcetype = Mellanox-Switch</small>

Patterns can be viewed not on real time and you can create alert on most repeatable events.

11.3 Appendix: Show Commands Not Supported By JSON API

Configuration Management
show configuration text files *
show files debug-dump *
show files stats *
Logging
show log
show log continuous
show log continuous matching *
show log continuous not matching *
show log debug
show log debug continuous
show log debug continuous matching *
show log debug continuous not matching *
show log debug files
show log debug files *
show log debug files * matching *
show log debug files * not matching *
show log debug matching *

show log debug not matching *
show log files
show log files *
show log files * matching *
show log files * not matching *
show log matching *
show log not matching *
Scheduled Jobs
show jobs
show jobs *
Subnet Manager (SM)
show ib sm log
show ib sm log continuous
show ib sm log continuous matching *
show ib sm log continuous not matching *
show ib sm log matching *
show ib sm log not matching *
User Management and Security
show users history
show users history username *
User Interfaces
show cli
show cli max-sessions
show cli num-sessions
show terminal

12 Document Revision History

Version 3.11.2224, May 2024

There are no changes to this version of the user manual. For further information on bug fixes and improvements, see the release notes of this software version.

Version 3.11.2202, February 2024

Updated:

- The section "[IP Table Filtering Default Rules](#)"

Version 3.11.2016, January 2024

There are no changes to this version of the user manual. For further information on bug fixes and improvements, see the release notes of this software version.

Version 3.11.2006, December 2023

Updated:

- [Subnet Prefix Checking](#) section
- "[system profile](#)" command notes

Version 3.11.2002, November 2023

Removed Signal Degradation Monitoring support

Added:

- IPv6 support for [tacacs-server host](#) command

Updated:

- "[show health-report](#)" command output
- "[show version](#)" command output and notification message
- "Ip filter" section under [Control Plane Policing \(CoPP\)](#)
- "[ip filter chain rule target | ipv6 filter chain rule target](#)" command
- "[Firewall Default Rules](#)" command

Version 3.11.10xx, July/September 2023

Updated:

- The subsection "[System Health Monitor Alerts Scenarios](#)"
- The section "[Cryptographic \(X.509, IPSec\) and Encryption](#)"

Version 3.10.60xx, April 2023

Added:

- A note to the command [ldap ssl](#)
- A note in the command [system profile](#)

Updated:

- The command [show interfaces ib transceiver](#)
- The section "[IB Router](#)"

Version 3.10.50xx, January 2023

There are no changes to this version of the user manual. For further information on bug fixes and improvements, see the release notes of this software version.

Version 3.10.41xx, November 2022

- Added note in the section "[Getting Started](#)"

Version 3.10.40xx, October 2022

Added:

- The ar-updn option to the command "[ib sm routing-engines](#)"

Removed:

- The command "ip l3" command
- Puppet Agent section

Version 3.10.31xx, August 2022

Updated:

- The command "[module-type](#)"

Version 3.10.30xx, July 2022

Added:

- The command "[ip filter reset-to-default-rules](#)"

Version 3.10.22xx, May 2022

There are no changes to this version of the user manual. For further information on bug fixes and improvements, see the release notes of this software version.

Version 3.10.21xx, April 2022

There are no changes to this version of the user manual. For further information on bug fixes and improvements, see the release notes of this software version.

Version 3.10.20xx, March 2022

Added:

- The command "[ldap nested-group-search](#)"
- The command "[ldap nested-group-depth](#)"
- The command "[ldap nested-group-count](#)"
- Note in the command "[system secure-mode enable](#)"

Updated:

- The command "[show ldap](#)"

Version 3.10.12xx, January 2022

There are no changes to this version of the user manual. For further information on bug fixes and improvements, see the release notes of this software version.

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. Neither NVIDIA Corporation nor any of its direct or indirect subsidiaries and affiliates (collectively: "NVIDIA") make any representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice. Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, and Mellanox are trademarks and/or registered trademarks of NVIDIA Corporation and/or



Mellanox Technologies Ltd. in the U.S. and in other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2024 NVIDIA Corporation & affiliates. All Rights Reserved.

