



Chassis Management

Table of contents

Chassis Management Commands

22

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

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

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.

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	A spine fan is not responsive on the switch system	Email, fan LED and system status LED	Check fan connectivity and replace it if required	"Fan <fan_number> in spine number

Alert Message	Scenario	Notification Indicator	Recovery Action	Recovery Message
<spine_number> is unresponsive		set red, log alert, SNMP		<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”

Alert Message	Scenario	Notification Indicator	Recovery Action	Recovery Message
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"
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"

Power Management

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.

Note

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

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.

Note

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
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

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>
```

- 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.

Monitoring Environmental Conditions

- 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
```

- Display measured voltage levels of power supplies.

```
switch (config) # show voltage
```

```
-----  
-----  
Module      Power Meter                               Reg          Expected  
Actual      Status  High   Low
```

Voltage		Range	Range	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      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
```

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

```
switch (config) # show power consumers
-----
-----
Module  Device           Sensor  Power      Voltage    Current    Status
Status
                [Watts] [Volts] [Amp]
-----
-----
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
```

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
```

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
```

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.

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

Note

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

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”.

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
```

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.

Note

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.

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.

Note

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
```

Note

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

System Reboot

Rebooting 1U Switches

To reboot a 1U switch system, run:

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

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
```

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	

Event Name	Description
NVE Decap fragmentation error	Fragmentation error during decapsulation

Chassis Management Commands

Chassis Management

clear counters

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

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		

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)

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 	

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

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.
-------	--

system profile

	<pre>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]</pre> <p>Sets the profile of the system to InfiniBand with various parameters</p>	
Syntax Description	ib-single-switch	<p>Enables InfiniBand switch profile All network interfaces link protocol set to InfiniBand</p>
	ib-no-adaptive-routing-single-switch	<p>Enables InfiniBand switch profile without adaptive routing capabilities All network interfaces link protocol set to InfiniBand with disabled adaptive routing</p>
	split-ready	<p>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.</p>
	ib-router	<p>Enables IB Routing capability on the system</p>
	num-of-swids	<p>Multiple switch IDs are configurable</p> <ul style="list-style-type: none"> • adaptive routing—enables adaptive routing • ib-router—enables IB router <p>Note: If num-of-swids is not defined then it is set to 1 by default.</p>
	no-adaptive-routing	<p>Disables adaptive routing</p>
	adaptive-routing-groups	<p>Sets adaptive routing groups.</p> <div style="background-color: #ffffcc; padding: 5px;"> <p> Note</p> </div>

		Allowed only when adaptive routing is enabled.
Default	The default system profile depends on the system.	
Configuration Mode	config	
History	3.1.0000 3.2.1100: Added "vpi-single-switch" option 3.3.4100: Added SX6036G3.3.4302Added system profile ib-no-adaptive-routing-single-switch 3.6.1002: Added system profile "ib num-of-swids" 3.6.6162: Added system profile "num of adaptive routing" 3.7.0020: Added system profile "ib split-ready" for 1U systems 3.8.1100: Updated syntax description for the parameter "adaptive-routing-groups" 3.9.0300: Added system profile "ib split-ready" for modular systems 3.9.2000: Updated note 3.10.6000: Updated note	
Example	switch (config) # system profile ib-single-switch	
Related Commands	port type show system profile show ports type	

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-single-switch” is the default profile for InfiniBand switches • The parameter “adaptive-routing-groups” is not available when “no-adaptive-routing” is set • 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. • Only NDR systems supports multi swid (num-of-swids can be larger than 1) , the maximum number of SWIDs is 8.
-------	---

usb eject

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

show asic-version

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

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	

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	

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	

show health-report

	<pre>show health-report Displays health report.</pre>
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	<pre>3.1.0000 3.3.0000: Output update 3.11.2000: Output update</pre>
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	

show inventory

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

Module Part Number Serial Number Asic Rev. HW Rev.	

CHASSIS 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	
Related Commands	
Notes	

show leds

	show leds [<module>] Displays the LED status of the switch system.
--	---

Syntax Description	module	Specifies the module whose LED status to display
Default	N/A	
Configuration Mode	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		

show memory

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

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</pre>

	<pre> ===== 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.

show power

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

show power consumers

	<pre> show power consumers Displays power consumption information. </pre>
--	---

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	

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	
-------	--

show resources

	show resources Displays system resources.
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	

show system capabilities

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

History	3.1.0000 3.3.0000: Added gateway support 3.6.1002: Updated example 3.7.0000: Updated example
Example	switch (config) # show system capabilities IB: Supported, L2, Adaptive Routing, Split Ready Max SM nodes: 648 IB Max licensed speed: EDR
Related Commands	show system profile
Notes	

show system manage inband-ib

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

show system profile

	show system profile Displays system profile.
--	---

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

show system type

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

show temperature

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

show version

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

	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
Related Commands	
Notes	

show version concise

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

show voltage

	show voltage Displays voltage level measurements on different sensors.
Syntax Description	N/A
Default	N/A

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

Chassis High Availability

chassis ha bip

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

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

chassis ha

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

chassis ha power enable other

	chassis ha power enable other no chassis ha power enable other Enables the other management card in the chassis. The no form of the command disables the other management card in the chassis.
Syntax Description	N/A
Default	The other management card is enabled

Configuration Mode	config
History	3.1.0000
Example	switch (config) # chassis ha power enable other
Related Commands	show chassis ha
Notes	This command is applicable only for modular switch systems.

show chassis ha

	show chassis ha 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	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
Related Commands	chassis ha
Notes	This command is applicable only for modular switch systems.

chassis ha bipv6

	ch assis 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	switch (config) # chassis ha bipv6 fdfd:fdfd:7:145::1033:47fd /64	
Related Commands	chassis ha	
Notes		

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. NVIDIA Corporation ("NVIDIA") makes no 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 NON-INFRINGEMENT, 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 and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

© Copyright 2024, NVIDIA. PDF Generated on 11/19/2024