



QM87xx 1U HDR 200Gb/s InfiniBand Switch Systems User Manual

QM8700, QM8790 1U HDR 200Gb/s InfiniBand Switch Systems User Manual

Table of Contents

Introduction.....	5
Speed and Switching Capabilities	6
Management Interfaces, PSUs and Fans	6
Features	7
Certifications	7
Installation.....	8
System Installation and Initialization	8
Safety Warnings	8
Air Flow.....	8
Package Contents	9
19” System Mounting Options	10
Fixed Rail Kit	10
Removing the System from the Rack.....	15
Telescopic Rail Kit	15
Removing the System from the Rack.....	19
Cable Installation	20
Splitter (Breakout) Cables and Adapters	21
Initial Power On	23
System Bring-Up of Managed Systems	24
Configuring Network Attributes.....	24
Remote Connection.....	27
FRU Replacements	28
Power Supply	28
Fans	29
Software Management	31
InfiniBand Subnet Manager	31
Upgrading Software (on Managed Systems)	31
Updating Firmware on Externally Managed Systems	32
Updating Firmware In-band (Typical).....	32
Interfaces	34
Data Interfaces.....	34
Speed	34

RS232 (Console)	34
Management	35
Micro USB	35
I ² C	35
Reset Button	36
LEDs	36
LED Notifications.....	36
System Status LED	37
Fan Status LED.....	37
Power Supply Status LEDs.....	38
Unit Identification LED	39
Port LEDs	40
Inventory Pull-out Tab.....	41
Troubleshooting	42
Specifications.....	43
Appendixes.....	44
Accessory and Replacement Parts	44
Thermal Threshold Definitions.....	44
Interface Specifications.....	45
QSFP Pin Description	45
RJ45 to DB9 Harness Pinout	47
Disassembly and Disposal	48
Disassembly Procedure	48
Disposal.....	48
Document Revision History	49

Relevant for Models: QM8700 and QM8790

This manual describes the installation and basic use of the Mellanox 1U HDR InfiniBand switch systems based on the Mellanox Quantum™ switch ASIC. This manual is intended for IT managers and system administrators.

Ordering Information

System Model	OPN	Legacy OPN	Description
QM8700	920-9B110-00-FH-0MD	MQM8700-HS2F	Mellanox Quantum™ HDR InfiniBand Switch, 40 QSFP56 ports, 2 Power Supplies (AC), x86 dual core, standard depth, P2C airflow, Rail Kit
	920-9B110-00-RH-0M0	MQM8700-HS2R	Mellanox Quantum™ HDR InfiniBand Switch, 40 QSFP56 ports, 2 Power Supplies (AC), x86 dual core, standard depth, C2P airflow, Rail Kit
QM8790	920-9B110-00-FH-0D0	MQM8790-HS2F	Mellanox Quantum™ HDR InfiniBand Switch, 40 QSFP56 ports, 2 Power Supplies (AC), unmanaged, standard depth, P2C airflow, Rail Kit
	920-9B110-00-RH-0D0	MQM8790-HS2R	Mellanox Quantum™ HDR InfiniBand Switch, 40 QSFP56 ports, 2 Power Supplies (AC), unmanaged, standard depth, C2P airflow, Rail Kit

Related Documentation

Document	Description
<i>InfiniBand Architecture Specification</i> Release 1.3.1, November 2, 2016 and Vol. 2, Release 1.4	The InfiniBand Trade Association (IBTA) InfiniBand® specification at https://www.infinibandta.org .
<i>MLNX-OS® User Manual</i>	This document contains information regarding the configuration and management of the MLNX-OS® software. See https://docs.nvidia.com/networking/category/mlnxos .
Hands-on workshops	https://academy.nvidia.com/en/infiniband-customized-training/
On-site/remote services	For any tailor-made service, contact: nbu-services-sales@nvidia.com .

Revision History

A list of the changes made to this document are provided in [Document Revision History](#).

Introduction

Mellanox QM8700/QM8790 switch systems provide the highest performing fabric solution in a 1U form factor by delivering up to 16Tb/s of non-blocking bandwidth with sub 130ns port-to-port latency. These switches deliver 7.2 billion packets-per-second (Bpps), or 390 million pps per port. These systems are the industry's most cost-effective building blocks for embedded systems and storage with a need for low port density systems. Whether looking at price-to-performance or energy-to-performance, these systems offer superior performance, power and space, reducing capital and operating expenses and providing the best return-on-investment.

Powerful servers combined with high-performance storage and applications that use increasingly complex computations are causing data bandwidth requirements to spiral upward. As servers are deployed with next generation processors, High-Performance Computing (HPC) environments and Enterprise Data Centers (EDC) need every last bit of bandwidth delivered with Mellanox's HDR InfiniBand systems.

Built with the latest Mellanox Quantum™ HDR 200Gb/s InfiniBand switch device, these standalone systems are an ideal choice for top-of-rack leaf connectivity or for building small to extremely large sized clusters. These systems enable efficient computing with features such as static routing, adaptive routing, and advanced congestion management. These features ensure the maximum effective fabric bandwidth by eliminating congestion. Mellanox Quantum™ is the world's smartest network switch, designed to enable in-network computing through Mellanox Scalable Hierarchical Aggregation Protocol (SHARP)™ technology. SHARP architecture enables the usage of all active data center devices to accelerate the communications frameworks, resulting in order of magnitude applications performance improvements.


Together with the Mellanox ConnectX®-6 adapter card, a Mellanox Quantum switch supports the HDR100 protocol. By utilizing two pairs of two lanes per port, a Mellanox Quantum switch can support up to 80 ports of 100G to create the highest density Top-of-Rack (ToR) switch available in the market. This is a perfect solution for double density racks with more than 40 servers per rack, as it also relieves small-medium deployments from the need to scale to 3-level fat-tree, and provides savings on power, latency and space.

The QM8700, dual-core x86 CPU, comes with an onboard subnet manager, enabling simple, out-of-the-box fabric bring-up for up to 2048 nodes. The QM8700 switch runs the same MLNX-OS® software package as Mellanox FDR and EDR products to deliver complete chassis management of the firmware, power supplies, fans and ports, and ensure interoperability with previous generation systems.

Mellanox's fixed-configuration systems can also be coupled with Mellanox's Unified Fabric Manager (UFM®) software for managing scale-out InfiniBand computing environments. UFM enables data center operators to efficiently provision, monitor and operate the modern data center fabric. UFM boosts application performance and ensures that the fabric is up and running at all times.

InfiniBand systems come as internally or externally managed. Internally managed systems include a CPU that runs the management software (MLNX-OS®) and management ports, which are used to transfer management traffic into the system. Externally managed systems come without the CPU and management ports, and are managed using firmware tools.

QM8700 Front View



QM8790 Front View



QM8700 and QM8790 Rear View



For additional airflow options, see [Airflow](#).

Speed and Switching Capabilities

The table below describes maximum throughput and interface speed per system model.

System Model	HDR 200Gb/s QSFP56 Interfaces	Max Throughput
QM8700	40	16Tb/s
QM8790	40	16Tb/s

Management Interfaces, PSUs and Fans

The table below lists the various management interfaces and available replacement parts per system model.

System Model	USB	MGT	I2C	Console	Replaceable PSU	Replaceable Fan
QM8700	Front (micro USB)	Front (1 port)	NA	Front	Yes	Yes
QM8790	NA	NA	Front (micro USB)	NA	Yes	Yes

Features

For a full feature list, please refer to the system's product brief. Go to <http://www.mellanox.com>. In the main menu, click on Products > InfiniBand/VPI Switch Systems, and select the desired product page.

Certifications

The list of certifications (such as EMC, Safety and others) per system for different regions of the world is located on the Mellanox website at http://www.mellanox.com/page/environmental_compliance.

Installation

System Installation and Initialization

Installation and initialization of the system require attention to the normal mechanical, power, and thermal precautions for rack-mounted equipment.



The rack mounting holes conform to the EIA-310 standard for 19-inch racks. Take precautions to guarantee proper ventilation in order to maintain good airflow at ambient temperature.



- Unless otherwise specified, Mellanox products are designed to work in an environmentally controlled data center with low levels of gaseous and dust (particulate) contamination.
- The operation environment should meet severity level G1 as per ISA 71.04 for gaseous contamination and ISO 14644-1 class 8 for cleanliness level.

The installation procedure for the system involves the following phases:

Step	Procedure	See
1	Follow the safety warnings	Safety Warnings
2	Pay attention to the air flow consideration within the system and rack	Air Flow
3	Make sure that none of the package contents is missing or damaged	Package Contents
4	Mount the system into a rack enclosure	19" System Mounting Options
5	Power on the system	Initial Power On
6	Perform system bring-up	System Bring-Up of Managed Systems
7	[Optional] FRU replacements	FRU Replacements

Safety Warnings

Prior to the installation, please review the [Safety Warnings](#). Note that some warnings may not apply to all models.

Air Flow

Mellanox systems are offered with two air flow patterns:


- Power (rear) side inlet to connector side outlet - marked with blue power supplies/fans FRUs' handles.

Air Flow Direction Marking - Power Side Inlet to Connector Side Outlet





- Connector (front) side inlet to power side outlet - marked with red power supplies/fans FRUs' handles.

Air Flow Direction Marking - Connector Side Inlet to Power Side Outlet



- All servers and systems in the same rack should be planned with the same airflow direction.
- All FRU components need to have the same air flow direction. A mismatch in the air flow will affect the heat dissipation.

The table below provides an air flow color legend and respective OPN designation.

Direction	Description and OPN Designation
	Connector side inlet to power side outlet. Red latches are placed on the power inlet side. OPN designation is "-R".
	Power side inlet to connector side outlet. Blue latches are placed on the power inlet side. OPN designation is "-F".

Package Contents

Before installing your new system, unpack it and check against the parts list below that all the parts have been sent. Check the parts for visible damage that may have occurred during shipping.

The QM8700 and QM8790 package content is as follows:

- 1 - System

- 1 - Rail kit
- 2 - Power cables Type C13-C14
- 1 x Harness: HAR000631 - Harness RS232 2M cable - DB9 to RJ-45 (only in QM8700)
- 2 - Cable retainers



If anything is damaged or missing, contact your sales representative at support@mellanox.com.

19” System Mounting Options

- By default, the systems are shipped with the static rail kit described in [Fixed Rail Kit](#).
- For the telescopic rail kit (purchased separately) installation instructions, see [Telescopic Rail Kit](#).

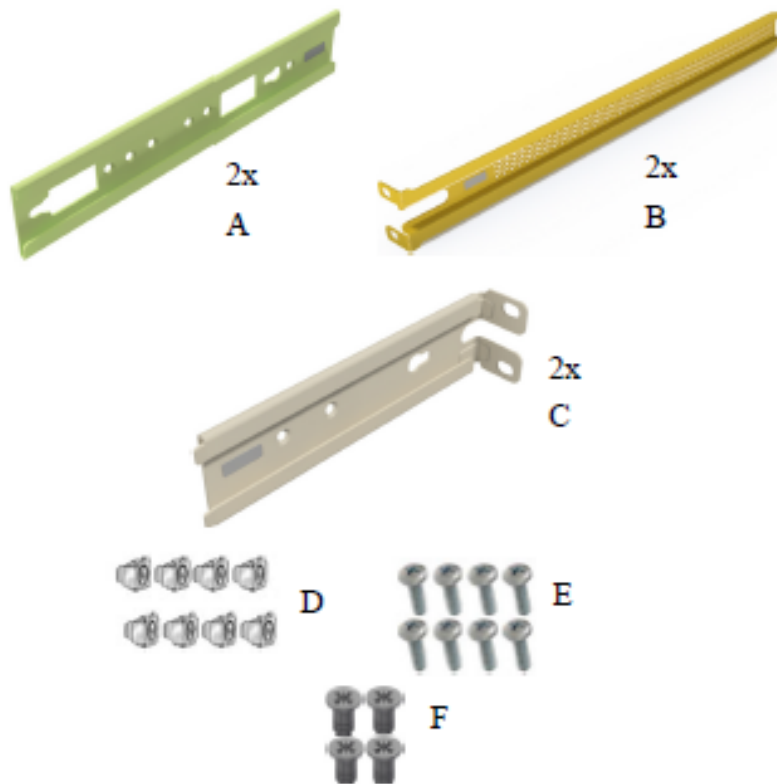
Fixed Rail Kit

Kit OPN	Legacy Kit OPN	Rack Size and Rack Depth Range
930-9BRKT-00JF-000	MTEF-KIT-C	430-800 mm

The following parts are included in the fixed rail kit (see figure below):

- 2x Rack mount rails (A)
- 2x Rack mount blades (B)
- 2x Rack mount ears (C)
- 8x M6 Standard cage nuts (D)
- 8x M6 Standard pan-head Phillips screws (E)
- 4x Flat Head Phillips 100 DEG 6-32X1/4" ST.ST PATCH 360 (F)

Rack Rail Kit Parts





Prerequisites:

Before mounting the system to the rack, select the way you wish to place the system. Pay attention to the airflow within the rack cooling, connector and cabling options.



While planning how to place the system, consider the two installation options shown in the figures below, and review the following points:

- Make sure the system air flow is compatible with your installation selection. It is important to keep the airflow within the rack in the same direction.
- Note that the part of the system to which you choose to attach the rails (the front panel direction, as demonstrated in Option 1 or the FRUs direction, as demonstrated in Option 2) will determine the system's adjustable side. The system's part to which the brackets are attached will be adjacent to the cabinet.
- The FRU side is extractable. Mounting the rack brackets inverted to the FRU side (Option 2) will allow you to slide the FRUs, in and out.

Short Racks (430-580 mm) Installation Options



Standard Racks (580-800 mm) Installation Options



⚠ In short racks, the system's ventilation openings should be framed by the designated windows in the rails, as shown below.

Short Racks (430-580 mm) Installation - Side View

Front side (ports):



Rear side (FRUs):



Standard Racks (580-800 mm) Installation - Side View

Front side (ports):



Rear side (FRUs):



To mount the system into the rack:


⚠ At least two people are required to safely mount the system in the rack.

⚠ The following steps include illustrations that show front side (ports) installation, yet all instructions apply to all installation options.

1. Attach the left and right rack mount rails (A) to the switch, by gently pushing the switch chassis' pins through the slider key holes, until locking occurs.
2. Secure the chassis in the rails by screwing 2 flat head Phillips screws (F) in the designated points with a torque of 1.5 ± 0.2 Nm.
Attaching the Rails to the Chassis



3. Attach the left and right rack mount ears (C) to the switch, by gently pushing the switch chassis' pins through the slider key holes, until locking occurs. Secure the system in the brackets by screwing the remaining 2 flat head Phillips screws (F) in the designated points with a torque of 1.5 ± 0.2 Nm.
Attaching the Brackets to the Chassis



4. Install 8 cage nuts (D) in the desired 1U slots of the rack: 4 cage nuts in the non-extractable side and 4 cage nuts in the extractable side.

Installing the Cage Nuts




⚠ While each rack U (unit) consists of three holes, the cage nut should be installed vertically with its ears engaging the top and bottom holes only.

While your installation partner is supporting the system's weight, perform the following steps:


5. Attach the two rack mount blades (B) to the back side (FRU side) of the rack by inserting four M6 screws (E) in the designated cage nuts. Do not tighten the screws yet.

Attaching the Brackets to the Rack



- Slide the switch with the rails (A) and ears (C) installed on it into the left and right rails (B) on the rack. Use four M6 screws (E) to fix the rack mount ears (C) to the rack. Do not tighten the screws yet.

Sliding the Blades in the Rails



- When fully inserted, fix the switch by tightening the 8 screws (E) inserted in Step 5 and Step 6 with a torque of 4.5 ± 0.5 .

Removing the System from the Rack

To remove a unit from the rack:

- Turn off the system and disconnect it from peripherals and from the electrical outlet.
While your installation partner is supporting the system's weight:
- Loosen the screws attaching the rack mount ears (C) to the rack. Do not remove them yet.
- Loosen the screws attaching the rack mount blades (B) to the rack, and pull the blades towards you, while your partner is holding the system.
- Extract the loosened screws from Step 2 and dismount the system from the rack.
- Remove the rails and brackets from the chassis by unscrewing 4 screws.

Telescopic Rail Kit

⚠ The telescopic rail kit is not included in the system's package, and can be purchased separately.

There are two installation kit options:

- Standard depth systems should be mounted using the standard rail kit.
- Short depth systems can be mounted using either of the rail kits.

Kit OPN	Legacy Kit OPN	Rack Size and Rack Depth Range
930-9BRKT-00JJ-000	MTEF-KIT-F	600-800 mm


The following parts are included in the rail kit package (see figure below):

- 1x Right side slider (A)
- 1x Left side slider (B)
- 2x Rear rail (C)
- 2x Front rail (D)
- 10x M6 Standard cage nuts^{1 2} (E)
- 10x M6 Standard pan-head Phillips screws¹ (F)
- 2x Phillips100 DEG F.H TYPE-I ST.ST 6-32 X 1/4 screw with around patch (G)
- 6x Flat head 100 DEG Phillips 4-40X3/16" ST.ST patch screws (H)

¹ Other threads are available by special order: M5, 10-32, 12-24

² G-type cage-nut is available by special order.

Rack Rail Kit Parts



Prerequisites

 The rails must be separated prior to the installation procedure.

To separate the rails:

1. Separate rail C from sliders A/B + D.
2. Extend the rail assembly by pulling the extension outwards (D).

Rails Separation



Before mounting the system to the rack, select the way you wish to place the system. Pay attention to the airflow within the rack cooling, connector and cabling options.

While planning how to place the system, review the following points:

- Make sure the system airflow is compatible with your installation selection. It is important to keep the airflow within the rack in the same direction.
- In case there are cables that cannot bend within the rack or in case more space is needed for cable bending radius, it is possible to recess the connector side or the FRU side by 3.15" (8cm) by optional placement of the system's rails.
- The FRU side is extractable. Mounting the sliding rail inverted to the system will allow you to slide the FRU side of the system, in and out.

To mount the system into the rack:

1. Install 10 cage nuts (E) into the desired 1U slot of the rack: 4 cage nuts in the non-extractable side and 6 cage nuts in the extractable side.

Installing the Cage Nuts




2. Mount both of the outer rails (C+D) into the rack (as illustrated below), and use 8 standard pan-head screws (F) to fix them to the rack. Do not tighten the screws yet.


Mounting the Outer Rails into the Rack



3. If cable accommodation is required, disassemble any of the inner rails from the brackets attached to them, by removing and scraping the connecting screws.
Disassembling the Inner Rails for Cable Accommodation




4. Route the power cable through either of the inner rails, and reassemble the brackets by screwing the 3 screws (per rail) provided with the rail-kit (H) with a torque of $0.7\pm 0.05\text{Nm}$.
Cable Accommodation




5. Secure the chassis in the inner rails screwing the 2 flat head Phillips screws (G) in the designated points with a torque of $1.5\pm 0.2\text{ Nm}$.

Securing the Chassis in the Inner Rails



6. Slide the switch into the rack by carefully pushing the inner rails into the outer rails installed on the rack.

Sliding the Switch into the Rack




7. When fully inserted, fix the switch by closing the remaining 2 screws in the middle and tightening the 8 screws inserted in Step 2 with a torque of 4.5 ± 0.5 Nm.

Removing the System from the Rack

To remove a unit from the rack:


1. Turn off the system and disconnect it from peripherals and from the electrical outlet.
2. Unscrew the two M6 screws securing the front of the inner rails' ears to the outer rails and to the rack.
3. Pull the unit out until braking is felt. For safety purposes, the locking mechanism will not allow a complete removal of the unit at this stage.

Pulling the Unit Outwards



4. Press on the locking spring (appears in red in the figure below) on both sides simultaneously, and continue pulling the unit towards you until it is fully removed.

Locking Mechanism



Cable Installation

All cables can be inserted or removed with the unit powered on.

To insert a cable, press the connector into the port receptacle until the connector is firmly seated. The LED indicator, corresponding to each data port, will light when the physical connection is established. When a logical connection is made, the relevant port LED will turn on.


To remove a cable, disengage the locks and slowly pull the connector away from the port receptacle. The LED indicator for that port will turn off when the cable is unseated.

For full cabling guidelines, ask your Mellanox representative for a copy of *Mellanox Cable Management Guidelines and FAQs Application Note*.

For more information about port LEDs, refer to [Port LEDs](#).

- ⚠ Do not force the cable into the cage with more than 40 newtons / 9.0 pounds / 4kg force. Greater insertion force may cause damage to the cable or to the cage.

Cable Orientation



Splitter (Breakout) Cables and Adapters



The breakout option is intended for users planning to run HDR100 using ConnectX-6 only.

The breakout cable is a unique capability, where a single physical quad-lane QSFP port is divided into 2 dual-lane ports. It maximizes flexibility by enabling end users to use a combination of dual-lane and quad-lane interfaces according to the specific requirements of their network. All system ports may be split into 2-lane ports. 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. For the systems splitting options, see “QM8700/QM8790 Splitting Options” below.

Sample Breakout Cable



- Splitting the interface deletes all configuration on that interface.
- This feature is available only for Quantum ASIC systems.
- In order to be able to use this feature, the system profile command must be activated with split-ready configuration (cross-reference to system profile command).

For more information on how to change the system's profile to allow Split-Ready configuration, how to change the module type to a split mode, and how to unsplit a split port, please refer to the "InfiniBand Switching" chapter in the latest *MLNX-OS® User Manual*.

QM8700/QM8790 Splitting Options



Split x2 All QSFP56 ports are splittable. Each port can be split to 2xQSFP56 (HDR100) ports. There are no blocking requirements.

Logical Port Numbering Schematic

Two profiles can be selected for the QM87x0 HDR switch systems. The first one defines the system as a pure 40-port HDR200 switch. The other profile permits any or all QSFP ports to be split into two 2X (HDR100) ports.

The following diagrams attempt to show how the logical ports map onto the physical QSFP ports, as viewed by the IB tools (e.g. ibnetdiscover):

Switch Profile: Non-Splittable (Suitable for L2/Spine Switches)

Physical Port #	1	3	5	7	9	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39
Logical 4X Port #	1	3	5	7	9	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39
Logical 4X Port #	2	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40
Physical Port #	2	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40



The IB tools report 41 logical ports. Port 41 is an internal port used for the SHARP Aggregation Node when SHARP is enabled.

Switch Profile: Splittable

Physical Port #	1	3	5	7	9	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39
Logical 2X Port #	1	5	9	13	17	21	25	29	33	37	41	45	49	53	57	61	65	69	73	77
Logical 2X Port #	2	6	10	14	18	22	26	30	34	38	42	46	50	54	58	62	66	70	74	78
Logical 2X Port #	3	7	11	15	19	23	27	31	35	39	43	47	51	55	59	63	67	71	75	79
Logical 2X Port #	4	8	12	16	20	24	28	32	36	40	44	48	52	56	60	64	68	72	76	80
Physical Port #	2	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40



Note: MLNX-OS will use Ethernet port numbering for split QSFP ports. For example, 2X ports 25 and 26 above will be named '1/13/1' and '1/13/2'.



Note: The IB tools will report 81 logical ports. Port 81 is an internal port used for the SHARP Aggregation Node when SHARP is enabled.

When the user wishes to keep a 4X port, rather than splitting it, from the IB tools view, the 4X port receives the odd port number, and the even-numbered port appears as disconnected. For example,

if physical Port 13 is not split, in MLNX-OS it will be referred to as '13', and the following scheme will apply:

Physical Port #	1	3	5	7	9	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39
Logical 2X Port #	1	5	9	13	17	21	25	29	33	37	41	45	49	53	57	61	65	69	73	77
Logical 2X Port #	2	6	10	14	18	22	26	30	34	38	42	46	50	54	58	62	66	70	74	78
Logical 2X Port #	3	7	11	15	19	23	27	31	35	39	43	47	51	55	59	63	67	71	75	79
Logical 2X Port #	4	8	12	16	20	24	28	32	36	40	44	48	52	56	60	64	68	72	76	80
Physical Port #	2	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40

⚠ MLNX-OS will refer to this 4X port as '1/13'.

Initial Power On

Each system's input voltage is specified in the [Specifications](#) chapter.

The power cords should be standard 3-wire AC power cords including a safety ground and rated for 15A or higher.


⚠ The system platform will automatically power on when AC power is applied. There is no power system. Check all boards, power supplies, and fan tray modules for proper insertion before plugging in a power cable.

1. Plug in the first power cable.
2. Plug in the second power cable.
3. Wait for the system upload process.

⚠ It may take up to five minutes to turn on the system. If the System Status LED shows amber after five minutes, unplug the system and call your Mellanox representative for assistance.

4. Check the System Status LEDs and confirm that all of the LEDs show status lights consistent with normal operation (initially flashing, and then moving to a steady color) as shown below. For more information, refer to [LED Notifications](#).

System Status LEDs 5 Minutes After Power On



⚠ After inserting a power cable and confirming the green System Status LED light is on, make sure that the Fan Status LED shows green. If the Fan Status LED is not green, unplug the power connection and check that the fan module is inserted properly and that the mating connector of the fan unit is free of any dirt and/or obstacles. If no obstacles were found and the problem persists, call your Mellanox representative for assistance.

Two Power Inlets - Electric Caution Notifications:



- Risk of electric shock and energy hazard. The two power supply units are independent. Disconnect all power supplies to ensure a powered down state inside of the switch platform.
- ACHTUNG Gefahr des elektrischen Schocks. Entfernen des Netzsteckers elnes Netzteils spannungsfrei. Um alle Einheiten spannungsfrei zu machen sind die Netzstecker aller Netzteile zu entfernen.
- ATTENTION Risque de choc et de danger e'lectriques. Le de'branchment d'une seule alimentation stabilise'e ne de'branch uniquement qu'un module "Alimentation Stabilise'e". Pour isoler completement le module en cause, Il faut de'brancher toutes les alimentations stabilise'es.
- 電擊與能源危害的危險。所有 PSU 均各自獨立。將所有電源供應器斷電，確保交換器平台內部在電源關閉狀態。

System Bring-Up of Managed Systems



The bring-up procedures described in this section do not apply to unmanaged/externally managed systems. Such systems are ready for operation after power-on.

In order to query the system, perform firmware upgrade or other firmware operation. Refer to the latest NVIDIA Firmware tools (MFT) located on <https://network.nvidia.com/products/adapter-software/firmware-tools/>.

In order to obtain the firmware version of the externally managed system:

1. Run the following command from a host:

```
# flint -d <device> q
```

2. Compare the results of this command with the latest version for your system posted on <https://network.nvidia.com/products/adapter-software/firmware-tools/>.
3. If the current version is not the latest version, follow the directions in the *MFT User Manual* to burn the new firmware.

Configuring Network Attributes

The procedures described in this chapter assume that you have already installed and powered-on the system according to the instructions in this document. Since the system comes with a pre-configured DHCP, you may find the explanation in [Disable Dynamic Host Configuration Protocol \(DHCP\)](#) sufficient. In case manual configuration is required, please refer to the instructions in [Manual Host Configuration](#).

Disable Dynamic Host Configuration Protocol (DHCP)

DHCP is used for automatic retrieval of management IP addresses.

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. In such a case, the serial connection should be used.





```
<localhost># ssh admin@<ip-address>

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

Manual Host Configuration

To perform initial configuration of the system:

Step 1. Connect a host PC to the Console RJ45 () port of the system, using the supplied harness cable (DB9 to RJ45). Make sure to connect to the Console RJ45 port and not to the (Ethernet) MGT () port.

Step 2. Configure a serial terminal program (for example, HyperTerminal, minicom, or Tera Term) on your host PC with the settings described in the table below. Once you perform that, you should get the CLI prompt of the system.

Serial Terminal Program Configuration

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

Step 3. Login as admin and use admin as password. On the first login, the MLNX-OS configuration wizard will start.

Step 4. To configure network attributes and other initial parameters to the system, follow the configuration wizard as shown in the Configuration Wizard Session table below.

Configuration Wizard Session

Wizard Session Display	Comments
Mellanox configuration wizard Do you want to use the wizard for initial configuration? yes	You must perform this configuration the first time you operate the system or after resetting the system. Type 'y' and then press <Enter>.
Step 1: Hostname? [switch]	If you wish to accept the default hostname, press <Enter>. Otherwise, type a different hostname and press <Enter>.
Step 2: Use DHCP on mgmt0 interface? [no] yes	Perform this step to obtain an IP address for the system. (mgmt0 is the management port of the system). 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 'no' (no Zeroconf), you must enter a static IP, and the session will continue.
Step 3: Enable IPv6? [yes]	The management interface will be able to use IPv6 addresses. If you enter "no" (no IPv6), you will automatically be referred to Step 6.
Step 4: Enable IPv6 auto-config (SLAAC) on mgmt0 interface? [no]	This turns on auto-configuration of the IPv6 addresses. This is unsuitable for DHCPv6.
Step 5: Enable DHCPv6 on mgmt0 interface? [no]	To enable DHCPv6 on the MGMT0 interface.
Step 6: Admin password (Press <Enter> to leave unchanged)? <new_password> Step 6: Confirm admin password? <new_password>	To avoid illegal access to the machine, please type a password and then press <Enter>. Then confirm the password by re-entering it. Note that password characters are not printed.
You have entered the following information: <A summary of the configuration is now displayed.> To change an answer, enter the step number to return to or hit <enter> to save changes and exit. Choice: <Enter> Configuration changes saved.	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. <i>Note: To re-run the configuration wizard, run the command "configuration jump-start" in Config mode.</i>

The table below shows an example of static IP configuration for mgmt0 interface.

Configuration Wizard Session - Static IP Configuration

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

Step 1: Hostname? []
Step 2: Use DHCP on mgmt0 interface? [yes] no
Step 3: Use zeroconf on mgmt0 interface? [no]
Step 4: Primary IP address? [for example 192.168.10.4] 10.10.10.10
Mask length may not be zero if address is not zero (interface eth0)
Step 5: Netmask? [0.0.0.0] 255.255.255.0
```

```
Step 6: Default gateway? [for example 192.168.10.1] 10.10.10.255
Step 7: Primary DNS server?
Step 8: Domain name?
Step 9: Enable IPv6? [yes]
Step 10: Enable IPv6 autoconfig (SLAAC) on mgmt0 interface? [no]
Step 11: Admin password (Enter to leave 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...

Step 5. Before attempting a remote (for example, SSH) connection to the system, check the mgmt0 interface configuration. Specifically, verify the existence of an IP address. To check the current mgmt0 configuration, enter the following command:

```
switch (config) # show interfaces mgmt0
Interface mgmt0 status:
  Comment:
  Admin up:          yes
  Link up:           yes
  DHCP running:     yes
  IP address:        10.209.28.50
  Netmask:           255.255.255.0
  IPv6 enabled:      yes
  Autoconf enabled: no
  Autoconf route:   yes
  Autoconf privacy: no
  DHCPv6 running:   no
  IPv6 addresses:    1
  IPv6 address:      fe80::202:c9ff:fe63:b55a/64
  Speed:             1000Mb/s (auto)
  Duplex:            full (auto)
  Interface type:    ethernet
  Interface source: physical
  MTU:               1500
  HW address:        00:02:C9:63:B5:5A

  RX bytes:          968810197      TX bytes:          1172590194
  RX packets:        10982099      TX packets:        10921755
  RX mcast packets: 0          TX discards:       0
  RX discards:       0          TX errors:         0
  RX errors:         0          TX overruns:       0
  RX overruns:       0          TX carrier:        0
  RX frame:          0          TX collisions:     0
                                TX queue len:      1000

switch (config) #
```

Step 6. Check the software version embedded in your system, using the command 'show version'. Compare this version to the latest version that can be retrieved from Mellanox support site. To upgrade software, please refer to the *MLNX-OS User Manual*.

Remote Connection

Once the network attributes are set, you can access the CLI via SSH or the WebUI via HTTP/ HTTPS.

To access the CLI, perform the following steps:

1. Set up an Ethernet connection between the system and a local network machine using a standard RJ45 connector.
2. Start a remote secured shell (SSH) using the command: `ssh -l <username> <IP_address>`

```
# ssh -l <username> <ip_address>
Mellanox MLNX-OS Switch Management
Password:
```

3. Login as admin (default username is *admin*, password is *admin*).
4. Once you get the CLI prompt, you are ready to use the system.

For additional information about MLNX-OS, refer to the *MLNX-OS User Manual* located on <https://docs.nvidia.com/networking/category/mlnxos>.

FRU Replacements

Power Supply

Mellanox systems are equipped with two replaceable power supply units work in a redundant configuration. Either unit may be extracted without bringing down the system.


❗ Make sure that the power supply unit that you are NOT replacing is showing all green, for both the power supply unit and System Status LEDs.

❗ Power supply units have directional air flows similar to the fan module. The fan module airflow must coincide with the airflow of all of the power supply units. If the power supply unit airflow direction is different from the fan module airflow direction, the system's internal temperature will be affected. For power supply unit air flow direction, refer to [Air Flow](#).

To extract a power supply unit:

1. Remove the power cord from the power supply unit.
2. Grasping the handle with your hand, push the latch release with your thumb while pulling the handle outward. As the power supply unit unseats, the power supply unit status LEDs will turn off.

3. Remove the power supply unit.
- PS Unit Pulled Out



To insert a power supply unit:

1. Make sure the mating connector of the new unit is free of any dirt and/or obstacles.

⚠ Do not attempt to insert a power supply unit with a power cord connected to it.

2. Insert the power supply unit by sliding it into the opening, until a slight resistance is felt.
3. Continue pressing the power supply unit until it seats completely. The latch will snap into place, confirming the proper installation.
4. Insert the power cord into the supply connector.
5. Insert the other end of the power cord into an outlet of the correct voltage.

⚠ The green power supply unit indicator should light. If it does not, repeat the whole procedure to extract the power supply unit and re-insert it.

Fans


The system can fully operate if one fan FRU is dysfunctional. Failure of more than one fan is not supported.

⚠ Make sure that the fans have the air flow that matches the model number. An air flow opposite to the system design will cause the system to operate at a higher (less than optimal) temperature. For power supply unit air flow direction, refer to [Air Flow](#).

To extract a fan unit:

1. Grasping the handle with your right hand, push the latch release with your thumb while pulling the handle outward. As the fan unit unseats, the fan unit status LEDs will turn off.
2. Remove the fan unit.

Fan Module Latches




To remove or replace a fan unit, gently pull out its black handle while pushing the latch release with your thumb.

To insert a fan unit:

1. Make sure the mating connector of the new unit is free of any dirt and/or obstacles.

2. Insert the fan unit by sliding it into the opening until slight resistance is felt. Continue pressing the fan unit until it seats completely.

 The green Fan Status LED should light. If not, extract the fan unit and reinsert it. After two unsuccessful attempts to install the fan unit, power off the system before attempting any system debug.

Software Management


The systems include an embedded management CPU card that runs MLNX-OS® management software.

The MLNX-OS systems management package and related documentation can be downloaded at <https://docs.nvidia.com/networking/category/mlnxos>.

InfiniBand Subnet Manager

The InfiniBand Subnet Manager (SM) is a centralized entity running in the system. The SM applies network traffic related configurations such as QoS, routing, partitioning to the fabric devices. You can view and configure the Subnet Manager parameters via the CLI/WebUI. Each subnet needs one subnet manager to discover, activate and manage the subnet.

Each network requires a Subnet Manager to be running in either the system itself (system based) or on one of the nodes which is connected to the fabric (host based).

 No more than two subnet managers are recommended for any single fabric.

The InfiniBand Subnet Manager running on the system supports up to 2048 nodes. If the fabric includes more than 2048 nodes, you may need to purchase NVIDIA's Unified Fabric Manager (UFM®) software package.

Each subnet needs one subnet manager to discover, activate and manage the subnet.

Each network requires a Subnet Manager to be running in either the system itself (system based) or on one of the nodes which is connected to the fabric (host based).

The subnet manager (OpenSM) assigns Local IDentifiers (LIDs) to each port connected to the fabric, and develops a routing table based on the assigned LIDs.

A typical installation using the OFED package will run the OpenSM subnet manager at system start up after the drivers are loaded. This automatic OpenSM is resident in memory, and sweeps the fabric approximately every 5 seconds for new adapters to add to the subnet routing tables.

Upgrading Software (on Managed Systems)

Software and firmware updates are available from the NVIDIA Support website. Check that your current revision is the same one that is on the NVIDIA website. If not upgrade your software. Copy the update to a known location on a remote server within the user's LAN.

Use the CLI or the GUI in order to perform software upgrades. For further information please refer to the Upgrading MLNX-OS® Software section in the [MLNX-OS Software User Manual](#).

Be sure to read and follow all of the instructions regarding the updating of the software on your system.

Managed systems do not require Firmware updating. Firmware updating is done through the MLNX-OS management software. The system comes standard with a management software module for system management called NVIDIA Operating System (MLNX-OS). MLNX-OS® is installed on all Quantum based managed systems. MLNX-OS® includes a CLI, WebUI, SNMP, system management software and IB management software (OpenSM).

Updating Firmware on Externally Managed Systems

There are two methods to update system firmware:

- (Typical) In-band via a switch network port across a cable connecting the server to the switch port.
- (Non-typical) Via the I²C port of the switch using a NVIDIA MTUSB-1 device connecting to a server's USB port on the one end and to the I²C port of the switch on the other.

Firmware updates should normally be conducted in-band. The use of the MTUSB-1 device is intended for cases of debug or firmware corruption and should be conducted by NVIDIA Fields or Support engineers, or by trained users at the customer's site.

Both types of updates require the installation of NVIDIA Firmware Tools (MFT) package. The MFT package and user manual are available for download under <https://network.nvidia.com/products/adapter-software/firmware-tools/>. Please select the package that suits your operating system.

In order to obtain information regarding the externally managed system, you must download the NVIDIA MFT tools from <https://network.nvidia.com/products/adapter-software/firmware-tools/>.

Select and download the release that matches your system. Follow the instructions in the User Manual <https://docs.nvidia.com/networking/category/mft> to get the tools.

Updating Firmware In-band (Typical)

Check the currently programmed firmware on the system and compare it to the latest firmware available under <https://network.nvidia.com/support/firmware/firmware-downloads/> (check under Quantum Switch Systems).

In order to obtain the firmware version of the externally managed system:

1. Obtain the LID of the target system. The following instructions use one of the utilities provided by the installed MFT package. (Other methods are described in the *MFT User Manual*) by performing the following:
 - a. Mark the GUID printed on the inventory pull-out tab of the system.
 - b. Run the command `ibnetdiscover` and search for the row starting with the word "Switch" and indicates the GUID of the system.
 - c. Mark the displayed LID on that row (a decimal number).
2. Run the following command from a host:

```
# flint -d <device> q#
```

3. Compare the results of this command with the latest version for your system posted on <https://network.nvidia.com/support/firmware/firmware-downloads/> (select the Quantum System page).
4. If the current version is not the latest version, follow the directions in the MFT User manual to burn the new firmware inband.

For further information, please refer to *MFT User Manual* at <https://docs.nvidia.com/networking/category/mft>.

Interfaces

The systems support the following interfaces:

- Data interfaces - InfiniBand
- 10/100/1000Mb Ethernet management interface (RJ45)*
- USB port (uUSB connector)
- RS232 Console port (RJ45)*
- I²C interface*
- Reset button
- Status and Port LEDs

*This interface is not found in externally managed systems.

In order to review the full configuration options matrix, refer to [Management Interfaces, PSUs and Fans](#).

Data Interfaces

The data interfaces use QSFP56 connectors. The full list of interfaces per system is provided in [Speed and Switching Capabilities](#).

Each QSFP56 port can be connected with QSFP56 cable or connector for 40/56/100/200 Gb/s. The system offers Class 7 (5W) QSFP56 transceivers support.


Speed

InfiniBand speed is auto-adjusted by the InfiniBand protocol. Mellanox systems support QDR/FDR/EDR/HDR InfiniBand.


- FDR is an InfiniBand data rate, where each lane of a 4X port runs a bit rate of 14.0625Gb/s with 64b/66b encoding, resulting in an effective bandwidth of 56.25Gb/s.
- EDR is an InfiniBand data rate, where each lane of a 4X port runs a bit rate of 25Gb/s with 64b/66b encoding, resulting in an effective bandwidth of 100Gb/s.
- HDR is an InfiniBand data rate, where each lane of a 4X port runs a bit rate of 50Gb/s with 64b/66b encoding, resulting in an effective bandwidth of 200Gb/s.


RS232 (Console)

The Console port labeled “**IOIOI**” is an RS232 serial port on the back side of the chassis that is used for initial configuration and debugging. Upon first installation of the system, you need to connect a PC to this interface and configure network parameters for remote connections. Refer to [Configuring Network Attributes](#) to view the full procedure.

 This interface is not found in externally managed systems.

Management


The Management RJ45 Ethernet ports labeled “” provide access for remote management. The management ports are configured with auto-negotiation capabilities by default (100MbE to 1000GbE). The management ports’ network attributes (such as IP Address) need to be pre-configured via the RS232 serial console port or by DHCP before use. Refer to [Configuring Network Attributes](#) to view the full procedure.

 This interface is not found in externally managed systems.


 Make sure you use only FCC compliant Ethernet cables.


Micro USB


The USB interface is USB 2.0 compliant and can be used by MLNX-OS software to connect to an external disk for software upgrade or file management. The connector comes in a standard micro USB shape. To view the full matrix of micro USB configuration options, refer to [Management Interfaces, PSUs and Fans](#).

- 
- USB 1.0 is not supported.
 - Do not use excessive force when inserting or extracting the USB disk to and from the connector.

I²C

The I²C connector is combined with the Console connector and is located on the front side of the system (the RJ45  connector). It can be used with the I²C DB9 to RJ45 splitting harness.

- 
- This interface is not found in managed systems. It is available in QM8790 systems only.
 - Apart from the initial configuration, I²C interface is made exclusively for debugging and troubleshooting. Only FAEs are authorized to connect through it.


 Only original Mellanox cables supplied with the switch package can be used to connect a switch system to the server.

Connecting any cable other than the Mellanox supplied console cable may cause an I²C hang. Using uncertified cables may damage the I²C interface.

Refer to the [Replacement Parts Ordering Numbers](#) appendix for harness details.

Reset Button

The reset button is located on the front side of the system next to the fan status LEDs. This reset button requires a tool to be pressed.

 Do not use a sharp pointed object such as a needle or a push pin for pressing the reset button. Use a flat object to push the reset button.

- To reset the system, push the reset button for less than 15 seconds.
- When using an Onyx (MLNX-OS) based system, keeping the reset button pressed for more than 15 seconds will reset the system and the “admin” password, this should allow you to enter without a password and set a new password for the user “admin”.





LEDs

See [LED Notifications](#).

LED Notifications


The system’s LEDs are an important tool for hardware event notification and troubleshooting.

LEDs Symbols

Symbol	Name	Description	Normal Conditions
	System Status LED	Shows the health of the system.	Green/Flashing green when booting
	Fan Status LED	Shows the health of the fans.	Green
	Power Supply Units LEDs	Shows the health of the power supply units.	Green
	Unit Identifier LED	Lights up on command through the CLI.	Off or blue when identifying a port

System Status LED

System Status LED - Front Side

Front Panel	Description
	The LED in the red oval shows the system's status.



⚠ It may take up to five minutes to turn on the system. If the System Status LED shows amber after five minutes, unplug the system and call your Mellanox representative for assistance.

System Status LED Assignments

LED Behavior	Description	Action Required
Solid Green	The system is up and running normally.	N/A
Flashing Green	The system is booting up. This assignment is valid on managed systems only.	Wait up to five minutes for the end of the booting process.
Solid Amber	Major error has occurred. For example, corrupted firmware, system is overheated, etc.	If the System Status LED shows amber five minutes after starting the system, unplug the system and call your Mellanox representative for assistance.

Fan Status LED

Fan Status LED - Front and Rear Sides

Front Panel	Description	Rear Panel
	Both of these LEDs in the red ovals show the fans' status.	

Fan Status Front LED Assignments

LED Behavior	Description	Action Required
Solid Green	All fans are up and running.	N/A
Solid Amber	Error, one or more fans are not operating properly.	The faulty FRUs should be replaced.

Fan Status Rear LED Assignments (One LED per Fan)

LED Behavior	Description	Action Required
Solid Green	A specific fan unit is operating.	N/A
Solid Amber	A specific fan unit is missing or not operating properly.	The fan unit should be replaced.



Risk of Electric Shock! With the fan module removed, power pins are accessible within the module cavity. Do not insert tools or body parts into the fan module cavity.

Power Supply Status LEDs

There are two power supply inlets in the system (for redundancy). The system can operate with only one power supply connected. Each power supply unit has a single 2 color LED that indicates the status of the unit.

Power Status LED



Rear Side Panel



Power Supply Unit Status Front LED Assignments

LED Behavior	Description	Action Required
Solid Green	All plugged (one or two) power supplies are running normally.	N/A

LED Behavior	Description	Action Required
Solid Amber	One or both of the power supplies are not operational or not powered up/ the power cord is disconnected.	Make sure the power cord is plugged in and active. If the problem resumes, the FRUs might be faulty, and should then be replaced.

The power supply status LEDs on the rear side of the system are located on the PSUs themselves. Each PSU has a single 2 color LED.

Power Supply Unit Status Rear LED Assignments

LED Behavior	Description	Action Required
Solid Green	All PS units are connected and running normally.	N/A
Flashing Green 1Hz	AC present / Only 12VSB on (PSU off) or PSU in Smart-on state.	Call your Mellanox representative for assistance.
Amber	AC cord unplugged or AC power lost while the second power supply still has AC input power.	Plug in the AC cord of the faulty PSU.
	PS failure (including voltage out of range and power cord disconnected).	Check voltage. If OK, call your Mellanox representative for assistance.
Flashing Amber	Power supply warning events where the power supply continues to operate; high temp, high power, high current, slow fan.	Call your Mellanox representative for assistance.
Off	No AC power to all power supplies.	Call your Mellanox representative for assistance.

Unit Identification LED

The UID LED is a debug feature, that the user can use to find a particular system within a cluster by turning on the UID blue LED.

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

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

To verify the LED status, run:

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

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


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

Port LEDs

By utilizing two pairs of two lanes per port, the systems can support up to 80 ports of 100G. You may switch between the two following states by pressing on the LED Splitting Control button:

- Displaying the link status of a single 4-lane port, or of the lower 2-lane split port (if a splitter cable is used). This is the default state.
- Displaying the link status of a single 4-lane port, or of the higher 2-lane split port (if a splitter cable is used).

Splitting Indication LEDs



Each time you press on the Lane Select Button, the Port LEDs display will switch to a different state, as follows:

Lane Select Button States

State	Indication LED 1	Indication LED 2	Port LEDs Display
1	On	Off	1st 2x
2	Off	On	2nd 2x

The port LEDs behavior indicates the ports' state, as follows:

Port LEDs in InfinBand System Mode

LED Behavior	Description	Action Required
Off	Link is down.	Check the cable.
Solid Green	Link is up with no traffic.	N/A
Flashing Green	Link is up with traffic.	N/A


LED Behavior	Description	Action Required
Solid Amber	Link is up.	Wait for the Logical link to raise. Check that the SM is up.
Flashing Amber	A problem with the link.	Check that the SM is up.

In InfiniBand system mode, the LED indicator, corresponding to each data port, will light orange when the physical connection is established (that is, when the unit is powered on and a cable is plugged into the port with the other end of the connector plugged into a functioning port). When a logical connection is made the LED will change to green. When data is being transferred the light will blink green.

Inventory Pull-out Tab

The system's inventory parameters (such as serial number, part number and GUID address) can be extracted from the inventory pull-out tab on the lower right side of the front panel.

Pull-out Tab



The images provided here are for illustration purposes only. They may not reflect the latest version of the product nor all available models.

Troubleshooting

Problem Indicator	Symptoms	Cause and Solution
LEDs	System Status LED is blinking for more than 5 minutes	<p>Cause: MLNX-OS software did not boot properly and only firmware is running.</p> <p>Solution: Connect to the system via the console port, and check the software status. You might need to contact an FAE if the MLNX-OS software did not load properly.</p>
	System Status LED is amber	<p>Cause:</p> <ul style="list-style-type: none"> • Critical system fault (CPU error, bad firmware) • Over temperature <p>Solution:</p> <ul style="list-style-type: none"> • Check environmental conditions (room temperature)
	Fan Status LED is amber	<p>Cause: Possible fan issue</p> <p>Solution:</p> <ul style="list-style-type: none"> • Check that the fan is fully inserted and nothing blocks the airflow • Replace the fan FRU if needed
	Front PSU Status LED is amber	<p>Cause: Possible PSU issue</p> <p>Solution:</p> <ul style="list-style-type: none"> • Check/replace the power cable • Replace the PSU if needed
	The activity LED does not light up (InfiniBand)	Make sure that there is an SM running in the fabric.
System boot failure	The last software upgrade failed on x86 based systems	<p>Solution:</p> <ul style="list-style-type: none"> • Connect the RS232 connector (CONSOLE) to a laptop. • Push the system's reset button. • Press the ArrowUp or ArrowDown key during the system boot. GRUB menu will appear. For example: <pre data-bbox="671 1503 1390 1809"> Default image: 'SX_X86_64 SX_3.4.0008 2014-11-10 20:07:51 x86_64' Press enter to boot this image, or any other key for boot menu Booting default image in 3 seconds. Boot Menu ----- 0: SX_X86_64 SX_3.4.0008 2014-11-10 20:07:51 x86_64 1: SX_X86_64 SX_3.4.0007 2014-10-23 17:27:34 x86_64 ----- Use the ArrowUp and ArrowDown 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: " </pre> <ul style="list-style-type: none"> • Select previous image to boot by pressing an arrow key and choosing the appropriate image.

Specifications

QM8700 and QM8790 Technical Specifications

Feature		Value
Mechanical	Size:	1.7" (H) x 17" (W) x 23.2" (D), 43.6mm (H) x 433.2mm (W) x 590.6mm (D)
	Mounting:	19" rack mount
	Weight:	1 PSU: 11.4 kg 2 PSUs: 12.488 kg
	Speed:	40, 56, 100, 200 Gb/s per port
	Connector cage:	40 QSFP56
Environmental	Temperature:	Operational: 0° to 40° C Non-Operational: -40° to 70° C
	Humidity:	Operational: 10%-85% non-condensing Non-Operational: 10%-90% non-condensing
	Altitude:	Operational: 3050m Non-Operational: 3050m
	Noise level:	71.9 dB(A)
Regulatory	Safety:	CB, cTUVus, CE, CU
	EMC:	EMC: CE, FCC, VCCI, ICES, RCM
	RoHS:	RoHS compliant
Power	Input Voltage:	100-127VAC 50/60Hz 4.5A; 200-240 50/60Hz 4.4A
	Global Power Consumption:	Full power specifications are provided in <i>NVIDIA QM87XX Full Technical Specifications</i> document, which is available in the customer portal following login.
Main Devices	CPU (in QM8700 only):	x86 ComEx Broadwell D-1508
	Switch:	NVIDIA Quantum™ IC
Throughput	Switching:	16Tb/s

*Measured with the following cables when running a stress test over all system ports:

MCP1650-H00AE30 - NVIDIA Passive Copper cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, 0.5m, black pulltab, 30AWG

MCP1650-H001E30 - NVIDIA Passive Copper cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, 1m, black pulltab, 30AWG

MCP1650-H002E26 - NVIDIA Passive Copper cable, IB HDR, up to 200Gb/s, QSFP56, LSZH, 2m, black pulltab, 26AWG

Appendixes

The document contains the following appendixes:

- [Accessory and Replacement Parts](#)
- [Thermal Threshold Definitions](#)
- [Interface Specifications](#)
- [Disassembly and Disposal](#)

Accessory and Replacement Parts

Ordering Part Numbers for Replacement Parts

OPN	Legacy OPN	Part Description
930-9BRKT-00JF-000	MTEF-KIT-C	Static rack installation kit for 200G 1U systems to be mounted into 430-800mm depth racks
930-9BRKT-00JJ-000	MTEF-KIT-F	Telescopic rack installation kit for 200G 1U systems to be mounted into 600-800mm depth racks
930-9BFAN-00IN-000	MTEF-FANF-C	200G 1U systems Fan Module w/ P2C airflow
930-9BFAN-00JO-000	MTEF-FANR-C	200G 1U systems Fan Module w/ C2P airflow
930-9BPSU-00JZ-000	MTEF-PSF-AC-C	200G 1U systems 1100W AC Power Supply w/ P2C airflow
930-9BPSU-00JG-000	MTEF-PSR-AC-C	200G 1U systems 1100W AC Power Supply w/ C2P airflow
HAR000631	-	Harness RS232 2M cable - DB9 to RJ-45 (for managed switches only)
ACC000501	-	Mellanox® 1U edge switch black power cord, 250V, 10A, 1830MM and C14 TO C13. Note: Can be purchased as a stand-alone product with PN ACC000501-BUY.

Thermal Threshold Definitions

Three thermal threshold definitions are measured by the Quantum™ ASICs, and impact the overall switch system operation state as follows:

- Warning - 105 °C: On managed systems only: When the ASIC device crosses the 100 °C threshold, a Warning Threshold message will be issued by the management software, indicating to system administration that the ASIC has crossed the Warning threshold. Note that this temperature threshold does not require nor lead to any action by hardware (such as switch shutdown).

- Critical - 120°C: When the ASIC device crosses this temperature, the switch firmware will automatically shut down the device.
- Emergency - 130°C: In case the firmware fails to shut down the ASIC device upon crossing its Critical threshold, the device will auto-shutdown upon crossing the Emergency (130°C) threshold.

Interface Specifications

QSFP Pin Description




QSFP Pin Description

Connector Pin Number	Pin Name	Signal Description
1	GND	Ground
2	Tx2n	Transmitter Inverted Data Input
3	Tx2p	Transmitter Non-Inverted Data Input
4	GND	Ground
5	Tx4n	Transmitter Inverted Data Input
6	Tx4p	Transmitter Non-Inverted Data Input
7	GND	Ground
8	Mod-SelL	Module Select

Connector Pin Number	Pin Name	Signal Description
9	ResetL	Module Reset
10	Vcc Rx	+3.3 V Power supply receiver
11	SCL	2-wire serial interface clock
12	SDA	2-wire serial interface data
13	GND	Ground
14	Rx3p	Receiver Non-Inverted Data Output
15	Rx3n	Receiver Inverted Data Output
16	GND	Ground
17	Rx1p	Receiver Non-Inverted Data Output
18	Rx1n	Receiver Inverted Data Output
19	GND	Ground
20	GND	Ground
21	Rx2n	Receiver Inverted Data Output 3
22	Rx2p	Receiver Non-Inverted Data Output 3
23	GND	Ground
24	Rx4n	Receiver Inverted Data Output 3
25	Rx4p	Receiver Non-Inverted Data Output 3
26	GND	Ground
27	ModPrsL	Module Present
28	IntL	Interrupt
29	Vcc Tx	+3.3 V Power supply transmitter
30	Vcc 1	+3.3 V Power Supply
31	LPMode	Low Power Mode
32	GND	Ground
33	Tx3p	Transmitter Non-Inverted Data Input
34	Tx3n	Transmitter Inverted Data Input
35	GND	Ground
36	Tx1p	Transmitter Non-Inverted Data Input
37	Tx1n	Transmitter Inverted Data Input
38	GND	Ground


QSFP Connector Male and Female Views



RJ45 to DB9 Harness Pinout

In order to connect a host PC to the Console RJ45 port of the system, a RS232 harness cable (DB9 to RJ45) is supplied.

RJ45 to DB9 Harness Pinout



RJ-45 Console and I²C interfaces are integrated in the same connector. Due to that, connecting any cable other than the Mellanox supplied console cable may cause an I²C hang.

Using uncertified cables may damage the I²C interface. Refer to the [Replacement Parts Ordering Numbers](#) appendix for harness details.

Disassembly and Disposal

Disassembly Procedure

To disassemble the system from the rack:

1. Unplug and remove all connectors.
2. Unplug all power cords.
3. Remove the ground wire.
4. Unscrew the center bolts from the side of the system with the bracket.

⚠ Support the weight of the system when you remove the screws so that the system does not fall.

5. Slide the system from the rack.
6. Remove the rail slides from the rack.
7. Remove the caged nuts.

Disposal

According to the WEEE Directive 2002/96/EC, all waste electrical and electronic equipment (EEE) should be collected separately and not disposed of with regular household waste. Dispose of this product and all of its parts in a responsible and environmentally friendly way.

Follow the instructions found at http://www.mellanox.com/page/dismantling_procedures for proper disassembly and disposal of the switch, according to the WEEE directive.



Document Revision History

Date	Revision	Description
July 31, 2022	2.0	Updated OPNs in: <ul style="list-style-type: none"> • Ordering Information • Fixed Rail Kit • Telescopic Rail Kit • Accessory and Replacement Parts
August 23, 2021	1.9	Updated Software Management .
March 29, 2020	1.8	Updated Thermal Threshold Definitions .
February 2020	1.7	<ol style="list-style-type: none"> 1. Updated I²C under Interfaces. 2. Added power consumption measuring details in Specifications.
January 2020	1.6	Added a warning to: <ol style="list-style-type: none"> 1. Interfaces under I²C section 2. Interface Specifications under RJ45 to DB9 Harness Pinout
October 2019	1.5	Global power consumption and noise level expanded in Specifications .
	1.4	Latency updated in Introduction .
July 2019	1.3	Removed preliminary from top of each page.
June 2019	1.2	Updated system description.
Jan. 2019	1.1	Migrated to online format; minor reorganization.
Nov. 2018	1.0	Initial release.

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

© 2022 NVIDIA Corporation & affiliates. All Rights Reserved.

