



NVIDIA Skyway-3 InfiniBand-to-Ethernet Gateway User Manual

Table of Contents

1	Introduction	6
1.1	Product Overview.....	6
1.1.1	NVIDIA Skyway-3 Highlights.....	6
1.2	Main System Components.....	7
1.2.1	Network Interface Cards.....	7
1.2.2	Power Supply Units.....	7
1.2.3	Fans	8
1.2.3.1	Power Supply Fans	8
1.2.3.2	Internal Fans	8
1.3	Package Contents.....	8
1.3.1	Rail Kit Package Contents	8
1.4	System Features.....	9
1.4.1	InfiniBand-to-Ethernet Gateway Operational Description	9
1.4.2	Load Balancing and High Availability Operational Description.....	10
1.5	Operating System.....	10
1.6	Certifications.....	10
2	System Layout and Interfaces.....	11
2.1	NVIDIA Skyway-3 Front and Rear Panel.....	11
2.2	Interfaces Detailed Description	12
2.2.1	Power-On LED.....	12
2.2.2	USB 3.0 Interfaces	12
2.2.3	PCIe Gen 5 Slots.....	13
2.2.4	LAN Interfaces	13
2.2.5	Redundant Power Module	13
2.2.6	Fans Modules.....	13
2.2.6.1	Power Supply Fans	13
2.2.6.2	Internal Fans	14
3	Hardware Installation	15
3.1	NVIDIA Skyway-3 Installation	15
3.1.1	Safety Warnings.....	15
3.1.2	System Requirements.....	16
3.1.2.1	Hardware Requirements	16
3.1.2.2	Airflow Requirements.....	16

3.1.2.3	Software Requirements	16
3.1.3	Unpacking the Package	16
3.1.3.1	Safety Precautions	16
3.1.3.2	System Package Contents.....	16
3.1.4	Rack Mounting	17
3.1.4.1	Installing the Server in a Rack	17
3.1.4.2	Slide Rail Installation	18
3.2	Cable Installation.....	23
3.2.1	Power Cable	23
3.2.2	ConnectX-6 Networking Cards Cables.....	23
3.2.2.1	Identifying Ethernet and InfiniBand/VPI Ports.....	23
3.3	System Connectivity	25
3.4	Initial Power-On	25
4	Configuring the Gateway for the First Time.....	26
4.1	Gateway Initialization.....	26
4.1.1	Rerunning the Wizard.....	31
4.2	Starting the Command Line Interface (CLI).....	31
5	Networkwide Deployment Guidelines	33
5.1	General Networkwide Guidelines.....	33
5.1.1	Ethernet Guidelines	33
5.1.2	InfiniBand Guidelines	34
5.1.2.1	Subnet Manager Configuration	34
5.2	Configuring IP Addresses and Routes.....	34
5.3	Configuring High Availability (HA)	36
5.3.1	Before Configuring HA	37
5.3.2	Configuring HA on Skyway Appliance.....	38
5.3.2.1	High Availability LAG/MLAG Setup	39
5.3.2.2	Skyway Connectivity to the Ethernet Using L2 Ethernet Switches	40
5.4	Configuring Partition Keys (PKEYs)	40
6	System Monitoring.....	42
6.1	Front Panel Monitoring Components.....	42
6.1.1	Power-On LED	42
6.2	Rear Panel LEDs.....	43
6.2.1	LAN Interfaces LEDs	43
6.2.1.1	LAN3/LAN4 Rear I/O LED Interface.....	43
6.2.1.2	LAN1/LAN 2 Rear I/O LED Interface	44

6.2.2	Power Module LED	44
6.2.3	Network Interface Cards LEDs.....	44
7	System Maintenance	46
7.1	Power Supply Units	46
7.1.1	Extracting and Inserting the Power Supply Unit.....	46
7.2	Slide Rail Kit	47
8	Troubleshooting	48
8.1	General Troubleshooting	48
9	Technical Specifications	49
9.1	MGA400-XS2 Specifications	49
9.2	System Dimensions.....	50
9.3	Thermal Threshold Definitions.....	50
10	Inventory Information	52
11	Field Replaceable Units	53
12	Revision History	54

About this Manual

This manual describes the installation and basic use of NVIDIA Skyway™-3 InfiniBand-to-Ethernet gateway.

Ordering Part Numbers

The table below provides the ordering part number (OPN) for the available NVIDIA Skyway-3 gateway.

NVIDIA SKU	Legacy OPN	Marketing Description
920-9B02D-00 RG-CF0	MGA400- XS2	NVIDIA Skyway-3 IB to Ethernet Gateway Appliance, 8x IB and 8x Ethernet ports, 2U server, 8x ConnectX-8 VPI dual-port QSFP112 NDR/400GbE adapters, 3 Power Supplies (AC)

Intended Audience

This manual is intended for IT managers and system administrators.

Related Documentation

Document Name	Description
NVIDIA MLNX-GW User Manual for NVIDIA Skyway	This document contains information regarding the configuration and management of the NVIDIA Skyway (MLNX-GW) software.

Revision History

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

1 Introduction

This is the user guide for the NVIDIA Skyway-3 InfiniBand-to-Ethernet gateway. This document contains the complete product overview, installation and initialization instructions, and product specifications.

1.1 Product Overview

NVIDIA Skyway-3 GA400 is an appliance-based InfiniBand-to-Ethernet gateway, enabling Ethernet storage or other Ethernet-based communications to access the InfiniBand datacenter, and vice versa. The solution, leveraging ConnectX's hardware-based forwarding of IP packets and standard IP-routing protocols, supports 400Gb/s NDR connectivity today, and is future-ready to support higher speeds.




1.1.1 NVIDIA Skyway-3 Highlights

Component	MGA400-XS2
Form Factor	2U rackmount: 19"
Weight	NVIDIA Skyway-3 gateway: 21kg The gateway with ACC and package: 32kg
PCIe Cards	8x NVIDIA® ConnectX®-8 dual-port network interface cards
InfiniBand/Ethernet Ports	8x InfiniBand ports 8x Ethernet ports
Connectivity Speed	InfiniBand: SDR/EDR/HDR100/HDR Ethernet: 25/50/100/200 Gb/s
Bandwidth	Up to 400Gb/s bi-directional per port
Power Supplies	3x AC power supplies

1.2 Main System Components

The NVIDIA Skyway-3 system populates eight ConnectX-6 InfiniBand/VPI adapter cards, fans, and two PSUs in the system's rear panel, as shown in the below figure.



1.2.1 Network Interface Cards



The following PCI speed/width degradation warnings for the Skyway-3 InfiniBand-to-Ethernet Gateway can be safely ignored:

- "Warning," "2024-01-01 01:00:00," "PCI Width Degradation Warning," "default(299) / Gateway: ibethgw:ib-gw / HCA-1/1," "IBPort," "PCI width degradation: enabled width is x16, active is x8 (depth 0 index 0 node 0)," 2.
- "Warning," "2024-01-01 01:00:00," "PCI Speed Degradation Warning," "default(299) / Gateway: ibethgw:ib-gw / HCA-1/1," "IBPort," "PCI speed degradation: enabled speed is Gen4, active is Gen3 (depth 0 index 0 node 0)," 2.

These warnings will also appear for other HCA cards: HCA-2/1, HCA-3/1, HCA-4/1, HCA-5/1, HCA-6/1, HCA-7/1, and HCA-8/1.

NVIDIA Skyway-3 is shipped populated with eight ConnectX-6 dual-port network interface cards (NICs) which enable the hardware-based forwarding of IP packets from the InfiniBand to Ethernet, and vice versa.

1.2.2 Power Supply Units

NVIDIA Skyway-3 is equipped with 2+1 redundant, load-sharing power supply units at the rear side of the system. The PSUs are housed in a 2U canister containing the power supplies. Each PSU has an extraction handle, PSU status LED, and a power socket.

For power supply unit LEDs operations, please refer to the [System Monitoring](#) section.

The system enables hot-swapping which enables components to be exchanged while the system is online without affecting operational integrity.



These power supply units can be removed from the system only if they are being replaced.

1.2.3 Fans

1.2.3.1 Power Supply Fans

NVIDIA Skyway-3 is equipped with one fan per power supply unit on the rear panel of the appliance.

1.2.3.2 Internal Fans

NVIDIA Skyway-3 is equipped with six internal fans for cooling the CPU and expansion cards. Under normal operation, the cooling fans operate at a constant speed. If the system module fails or one of the temperature thresholds is exceeded, the cooling fans automatically raise their rotation speeds to draw more airflow.

1.3 Package Contents

Check the package contents list to see that all the parts have been sent. Check the parts for visible damage that may have occurred during shipping. Please note that the product must be placed on an antistatic surface.

Category	Qty.	Item
Systems	1	NVIDIA Skyway-3 2U system
Slide Rail Kit	1	2U 36" slide kit pair for NVIDIA Skyway-3
Power Cables	2	250V 10A 1830MM C14 TO C13 power cable
	2	Cable retainers
Harness	1	Harness RS232 2M cable—DB9 to RJ-45 (do not connect to the COM port)
Documentation	1	Quick Installation Guide

1.3.1 Rail Kit Package Contents

Category	Qty.	Item
Slides	1	2 sets of slides
Screw M5* 15L	2	8 pcs

Management Interfaces, PSUs, and Fans

Processor System	Chipset	Intel 4209T, 2.2GHz, 11M, 8 Cores
	CPU Type	Dual Intel Xeon Silver 4516Y Scalable processor (up to 185W TDP)
	Memory Type	Supports DDR5 2133/2400/2666 MHz ECC-REG Modules
	Memory Size	16 x 32GB DDR5 6400MHz
	Memory Voltage	1.2V
	Error Detection	• Corrects single-bit errors • Detects double-bit errors (using ECC memory)
Rear I/O Panel	USB	2x USB 3.0
	RJ-45/LAN	3 x RJ-45 LAN ports: • 2x 1GbE • 1x 1GbE/IPMI-LAN
On-board Devices	EC	TE 8528E chip provides motherboard, RS-232, and hardware monitor functions
	BMC	Sharing with the LAN 1/3.
Expansion slots	PCI-Express	8x network interface cards
Cooling	Chassis Fan	6x 80cm high-speed fans
	PSU Fans	One fan per power supply unit
PC Health Monitoring	Voltage	Monitors for CPU Cores, +3.3 V, +5V, +12V, +5V standby, VBAT
	Temperature	Monitoring for CPU0 & CPU1 (PECI) Monitoring for system (HWM)
	Other Features (Case Open)	Chassis intrusion detection

1.4 System Features

For a full list of features, please refer to the system's product brief at www.nvidia.com/en-us/networking. In the main menu, click on PRODUCTS → INFINIBAND → GATEWAY & ROUTERS SYSTEMS → select the desired product page.

1.4.1 InfiniBand-to-Ethernet Gateway Operational Description

NVIDIA Skyway-3 GA100 is an appliance-based InfiniBand-to-Ethernet gateway, enabling Ethernet storage or other Ethernet-based communications to access the InfiniBand datacenter, and vice versa. The solution, leveraging ConnectX's hardware-based forwarding of IP packets and standard IP-routing protocols, supports 200Gb/s HDR connectivity, today, and is future-ready to support higher speeds.

NVIDIA Skyway-3 contains 8 ConnectX VPI dual-port adapter cards which enable the hardware-based forwarding of IP packets from InfiniBand to Ethernet, and vice versa. NVIDIA Skyway-3 also includes the NVIDIA Gateway Operating System, MLNX-GW, which manages the appliance and handles the high availability and load balancing between the ConnectX cards and gateway appliances.

A single NVIDIA Skyway-3 supports a maximum bandwidth of 1.6Tb/s, utilizing 16 ports, each of which reaches 100Gb/s traffic. In terms of connectivity, the InfiniBand ports can be connected to the InfiniBand network using HDR/HDR100 or EDR speeds, while the Ethernet ports can be connected to the Ethernet network using 200Gb/s or 100Gb/s.

1.4.2 Load Balancing and High Availability Operational Description

On the Ethernet side, the load balancing and high availability functions are achieved by leveraging Ethernet LAG (Link Aggregation). LACP (Link Aggregation Control Protocol) is used to establish the LAG and to verify connectivity. On the InfiniBand side, these functions are achieved by assuring that different flows go through different ConnectX HCAs, so that, in case a HCA drops, another HCA will continue passing its flows.

At initialization, 256 gateway GIDs are spread evenly among all InfiniBand ports of the gateway appliances. When an InfiniBand node initiates a traffic flow through the gateway, it first sends a broadcast ARP request with the default gateway IP address to determine the gateway's GID. All ConnectX cards receive the request, but only one sends the ARP response. The response is sent from the ConnectX cards that were assigned to handle the range of GIDs corresponding to the sending node's IP address. When the originating node receives the gateway GID, it sends a path query to the subnet manager (SM) to determine the gateway LID and the communication flow is performed as usual.

The dynamic assignment of the 256 gateway GIDs is the basic element of the load balancing and high availability operations. For any change in gateway configurations (e.g., a cable is dropped, an Ethernet link is disabled, or an appliance is powered off), the gateway GIDs get reassigned by MLNX-GW to be handled by other ConnectX cards. From the end-node perspective, nothing is changed. The same GID and LID remain, even when handled by a different ConnectX HCA.

1.5 Operating System

NVIDIA Skyway-3 includes the NVIDIA Gateway operating system, MLNX-GW, which manages the appliance and handles the high availability and load balancing between the ConnectX cards and between gateway appliances. For a detailed description of MLNX-GW, see please see [NVIDIA MLNX-GW User Manual for NVIDIA Skyway-3](#) or contact your NVIDIA representative.

1.6 Certifications

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

2 System Layout and Interfaces

The figures below show the front and rear sides of NVIDIA Skyway-3. Each numbered interface that is referenced in the figures is described in the following table with a link to detailed information.

2.1 NVIDIA Skyway-3 Front and Rear Panel



Item	Interface	Description
1	Power-On LEDs	Bezel LED—lights when the system is powered on
3	USB Interfaces	2 x USB 3.0 ports at the rear panel
4	PCIe Gen 5 Slots	8x ConnectX-8 VPI adapter cards
5	NIC ports	2x 1GbE Base-T ports

Item	Interface	Description
6	RJ-45 Port	This port is designed for secure local and remote server management control (BMC) and helps IT administrators deploy, update and monitor the device.
7	Mini Display Port	Mini display port video connector
8	Drive slots	1x E1.S NVMe drives. Enables you to install drives that are supported on your system
	Pull-Tab	A slide-out label panel containing system information (e.g., MAC address, serial number, etc.)
	Control Panel	Contains the system health, system ID, and status LED
9	Redundant Power Module	2+1 redundant 80 PLUS PLATINUM PSUs at the rear of the appliance
10	Power Supply Fans	One fan per power supply unit

2.2 Interfaces Detailed Description

2.2.1 Power-On LED

There is one I/O LED (green) on the front panel, to indicate if the system is powered.


- For Power-On LEDs definitions, please refer to [Power-On LEDs Specifications](#)

2.2.2 USB 3.0 Interfaces

Skyway-3 offers 2x USB 3.0 ports on the system's rear panel.

The USB interfaces are USB 3.0 compliant and can be used to provide the bandwidth up to 500MB/s to shorten the time for data transmission.

4x USB 3.0



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

2.2.3 PCIe Gen 5 Slots

Skyway-3 is shipped assembled with eight ConnectX-8 VPI cards that enable the hardware-based forwarding of IP packets from InfiniBand to Ethernet, and vice versa.

2.2.4 LAN Interfaces

Skyway-3 features two GbE LAN 1/2 (Intel i350 (2* 10/100/1000 Base-T GbE)) to eliminate bottlenecks in network data flow. The ports are accessible from the appliance's rear panel.

In addition, the appliance is equipped with ASPEED 2600 BMC chip that supports IPMI 2.0 (Intelligent Platform Management Interface 2.0) via RJ-45 port.

The recommended cabling when connecting LAN1 and LAN2 (for 10G LAN) is Cat. 6 (< 55m) or Cat. 6A (<100m) for which both cross and straight cables are supported.

- For LAN1 and LAN2 LED definitions, please refer to [LAN1 / LAN 2 Rear I/O LED Specifications](#)
- For LAN3 and LAN4 LEDs definitions, please refer to [LAN3 / LAN4 Rear I/O LED Specifications](#)

2.2.5 Redundant Power Module

NVIDIA Skyway-3 is equipped with 3 redundant power supply units at the rear of the appliance. The PSUs are housed in a 2U canister containing the power supplies. Each PSU has an extraction handle, PSU status LED, and a power socket.

- 80 PLUS Platinum 2+1 redundant power supply
 - 1000 W @ 100 ~ 127 V
 - 2000 W @ 200 ~ 240 V
-
- For the redundant power module LEDs definitions, please refer to [Power Module LED specifications](#)

2.2.6 Fans Modules

2.2.6.1 Power Supply Fans

NVIDIA Skyway-3 is equipped with one fan per power supply unit on the rear panel of the appliance.



PSU Fans

2.2.6.2 Internal Fans

NVIDIA Skyway-3 is equipped with six internal fans for cooling the CPU and expansion cards. Under normal operation, the cooling fans operate at a constant speed. If the system module fails, or one of the temperature thresholds are exceeded, the cooling fans automatically raise their rotation speeds to draw more airflow. For example, if there are two fans for CPU cooling and the temperature of CPU0 raises, only these two system fans will raise their fan speed and the other four fans for the cooling of expansion cards will operate at a constant speed. Below is a list of the different internal fan types.

- 6x 8cm heavy duty fans with optimal fan speed control
- All fans are equipped with tachometer status monitoring

3 Hardware Installation

Installation of the NVIDIA Skyway-3 gateway requires attention to the mechanical and power elements of the appliance and precautions must be taken for the rack-mounted equipment.

The system platform can be rack-mounted and is designed for installation in a standard 19" rack. The power side of the system includes two hot-swap power supply units (PSUs), and replaceable fan trays. There is one possible airflow direction. It is necessary to validate that the system airflow direction is compatible with the system, rack, and PSUs. The rear panel of the system has the QSFP28 ports, system LEDs, and management connection ports.

Use a rack capable of supporting the mechanical and environmental characteristics of a fully-populated platform.



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.

3.1 NVIDIA Skyway-3 Installation

The installation procedure of NVIDIA Skyway-3 involves the following steps.

Step	Procedure	Direct Link
1	Follow safety warning procedures.	Refer to Safety Warnings
2	Pay attention to the system considerations within the host chassis.	Refer to System Requirements
3	Follow the safety precautions	Refer to Safety Precautions
4	Unpack the package and confirm all required components are present.	Refer to Unpacking the Package
5	Mount the appliance in a rack enclosure.	Refer to Rack Mounting
6	Use the supplied cables to connect the system	Refer to Cable Installation
7	Power on the system.	Refer to Initial Power-On

3.1.1 Safety Warnings

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

3.1.2 System Requirements

3.1.2.1 Hardware Requirements



Unless otherwise specified, NVIDIA Networking products are designed to work in an environmentally controlled data center with low levels of gaseous and dust (particulate) contamination.

The operating environment should meet severity level G1 as per ISA 71.04 for gaseous contamination and ISO 14644-1 class 8 for cleanliness level.

3.1.2.2 Airflow Requirements

NVIDIA Skyway-3 appliance is offered with one airflow pattern: from the front panel to the rear panel.

Refer to the [Technical Specifications](#) section for airflow numbers.

3.1.2.3 Software Requirements

See [Operating Systems](#) section under the Introduction section.

3.1.3 Unpacking the Package


3.1.3.1 Safety Precautions

The NVIDIA Skyway-3 appliance is installed in systems that operate with voltages that can be lethal. Before opening the case of the system, observe the following precautions to avoid injury and prevent damage to system components.

- Remove any metallic objects from your hands and wrists.
- Make sure to use only insulated tools.
- Verify that the system is powered off and is unplugged.
- Place the ESD mat on the floor where working and put on the ESD strap. Make sure the ESD strap is touching your skin and that the other end is connected to a verified ground.

3.1.3.2 System Package Contents

Check the package contents list to see that all the parts have been sent. Check the parts for visible damage that may have occurred during shipping. Please note that the product must be placed on an antistatic surface. Please refer to [Package Contents](#).




The device comes with a key to the front panel. Note that although the key is provided, there is no reason to open the front panel nor to access the elements behind the front panel.

3.1.4 Rack Mounting

The NVIDIA Skyway-3 appliance can be mounted in a rack using the optional rack mounting kit. We strongly recommend that the minimum depth of cabinet is 1100mm.

3.1.4.1 Installing the Server in a Rack

Before mounting the NVIDIA Skyway-3 appliance in a rack, ensure that all internal components have been installed and that the unit has been fully tested. Both sides of the chassis ear must be assembled with screws (PN:1930005209) after the slide rail kit has been assembled.




3.1.4.2 Slide Rail Installation




Read prior to installation.


The server slides are developed for 1U or 2U applications of which system load does not exceed 75lbs. The slide length is 1041 ± 3.0 mm. The rear bracket is extendable to a max/min post-to-post distance of 670-1042 mm. The slide extension is 610.0 ± 3.0 mm.




Step 1: Remove inner member. Pull inner member out as in the illustration.




Step 2: Mount the inner member onto the chassis. Place the key slot on T stud, and push the inner member toward the back.




Step 3: Mount the cabinet member to the posts. Align the positioning pin to the desired complete U location, and pull the bracket forwards to lock it to the post. The bracket is locked to the post after a “click” sound is heard.



Step 4: Release the locking latch upward.



Step 5: Push the middle member forward to the rear of the slide.




Step 6: Install the chassis.

As shown, insert the inner member to the cabinet member. Make sure the ball retainer is in the open position. If the ball retainer is not on the front position, it might cause damage to the slides. After the inner member goes in, push up/down the disconnect lever to unlock the slides and keep


pushing the chassis to the fully-closed position.

Attention: ball retainer must be in the open position.



Push up / down the lever

Step 7: Screw the system in the cabinet.



Rack Mount Instructions (similar rack-mount instructions are included with the installation instructions):

- Elevated Operating Ambient—If installed in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment may be greater than the room ambient. Therefore, consideration should be given to installing the equipment in an environment compatible with the maximum ambient temperature (T_{ma}) specified by the manufacturer.
- Reduced Air Flow— Installation of the equipment in a rack should be such that the amount of airflow required for the safe operation of the equipment is not compromised.
- Mechanical Loading—Mounting of the equipment in the rack should be such that a hazardous condition is not achieved due to uneven mechanical loading.
- Circuit Overloading—Consideration should be given to the connection of the equipment to the supply circuit and the effect that overloading of the circuits

might have on overcurrent protection and supply wiring. Appropriate consideration of equipment nameplate ratings should be used when addressing this concern.

- Reliable Earthing—Reliable earthing of rack-mounted equipment should be maintained. Particular attention should be given to supply connections other than direct connections to the branch circuit (e.g., use of power strips).
- Please note that the handlebar and mounting ear must be installed after the slide rail kit has been installed completely.

3.2 Cable Installation

3.2.1 Power Cable

The NVIDIA Skyway-3 appliance is shipped with two power supply units. Each supply unit has a separate AC receptacle. The appliance accepts voltages of 100-127 VAC and 200-240 VAC for all possible power supply units. The power cords should be a standard 3-wire AC power cords, including a safety ground, and rated for 15A or higher. The power supplies deliver 2KW AC.


✗ After inserting a power cable and turning the appliance on, confirm the green system LED light is on.

✗ Do not hot swap the power supply if your appliance has only one power supply. Instead, power down the system to replace the power supply unit.

3.2.2 ConnectX-6 Networking Cards Cables

To obtain the list of supported NVIDIA Networking cables for the adapter cards, please refer to the Cables Reference Table at <http://www.mellanox.com/products/interconnect/cables-configurator.php>.

3.2.2.1 Identifying Ethernet and InfiniBand/VPI Ports



Networking Cable Installation

1. All cables can be inserted or removed with the unit powered on.
2. To insert a cable and press the connector into the port receptacle until the connector is firmly seated.
 - a. Support the weight of the cable before connecting the cable to the adapter card. Do this by using a cable holder or tying the cable to the rack.
 - b. Determine the correct orientation of the connector to the card before inserting the connector. Do not try and insert the connector upside down. This may damage the adapter card.
 - c. Insert the connector into the adapter card. Be careful to insert the connector straight into the cage. Do not apply any torque, up or down, to the connector cage in the adapter card.
 - d. Make sure that the connector locks in place.



When installing cables make sure that the latches engage.



Always install and remove cables by pushing or pulling the cable and connector in a straight line with the card.

3. After inserting a cable into a port, the green LED indicator will light 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). See [Network Interface Cards LEDs](#) under the Interfaces section.
4. After plugging in a cable, lock the connector using the latching mechanism particular to the cable vendor. When data is being transferred, the green LED will blink. See [Network Interface Cards LEDs](#) under the Interfaces section.
5. Care should be taken to not impede the air exhaust flow through the ventilation holes. Use cable lengths that allow for routing horizontally around to the side of the chassis before bending upward or downward in the rack.
6. To remove a cable, disengage the locks and slowly pull the connector away from the port receptacle. The LED indicator will turn off when the cable is unseated.

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 up 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 NVIDIA Networking representative for a copy of NVIDIA Cable Management Guidelines and FAQs Application Note.



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

3.3 System Connectivity

This section indicate how the ConnectX-6 IB ports should be physically connected to the IB switches. The below section describes the following topologies:

(e.g. with split-cables, providing 2x HDR100 connections from 2x ConnectX-6(s) to 1 Quantum port).

and what I am missing is like a "zoom in" into the actual Skyway-3 , show how to connect splitter cables from the IB switch to it and how to use regular HDR cables

3.4 Initial Power-On

The system's input voltage is specified in the [Technical Specifications](#) chapter. The power cords should be a 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 fans for proper insertion before plugging in a power cable.

Step 1. Plug in the first power cable.

Step 2. Plug in the second power cable.

Step 3. Wait for the System Status LED to turn green.



It may take up to five minutes to turn on the system. If the System Status LED is red after five minutes, unplug the system and call your NVIDIA Networking representative for assistance.

Step 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). For more information, refer to [System Monitoring](#).



After inserting a power cable and confirming the green System Status LED light is on, make sure that the Fan Status LED is 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 NVIDIA Networking representative for assistance.

4 Configuring the Gateway for the First Time

4.1 Gateway Initialization


To initialize the gateway, follow the steps below.

1. Enable remote access to serial console via IPMI.




Steps a through d, below, instruct how to find the MAC address for the IPMI port from inside the BIOS, even though the outside of the chassis is labeled with the IPMI port MAC address.

- a. Connect a VGA monitor and USB keyboard directly to the NVIDIA Skyway appliance.
- b. To enter the BIOS, reboot the NVIDIA Skyway appliance and press during bootup until the BIOS window pops up.
- c. Go to "Server Mgmt." tab → "BMC network configuration."
- d. The "Station IP address" is the address of the IPMI controller. DHCP may need to be configured in order to provide a lease for the MAC address. The NVIDIA Skyway appliance has 2 LAN ports on the back panel of the appliance that can be used for IPMI (in the figure below IPMI LAN2 is used).



"Lan channel 1" refers to IPMI LAN 1 and "Lan channel 2" refers to IPMI LAN 2.



- e. Use the following IPMI command to remote access serial console (user and password should be “admin” by default).

```
ipmitool -I lanplus -H <IPMI_CONTROLLER_IP> -U <user> -P <password> sol activate
```



The command should be run on a Linux console with the “ipmitool” application installed.

Example:

```
ipmitool -I lanplus -H 10.7.113.60 -U admin -P admin sol activate
```



Make sure to connect to the console SOL port of the gateway and not to the management port. Of the four ports, either of the outer ports (1st or 4th port—either to IPMI LAN1 or IPMI LAN 4, in the image above) can be selected to be the SOL port.



Once operating system boots, iKVM over HTML5 no longer shows any output. However, iKVM over HTML5 can be used for BIOS configurations at the very beginning of the system boot sequence right before the operating system boots.



By default, the VMC port (5900) is enabled, with no service (or VNC) running in the background. If needed, the port can be closed by disabling KVM using the IPMI tool as follows:

To disable KVM (and VNC):

```
ipmitool -I lanplus -H ${BMC_IP} -U ${username} -P ${password} raw 0x3e 0xc0 0x0c 0x01 0x00 0x04
```

To enable KVM:

```
ipmitool -I lanplus -H ${BMC_IP} -U ${username} -P ${password} raw 0x3e 0xc0 0x0c 0x01 0x01 0x04
```

2. Configure Console Redirection. This configuration allows to use remote IPMI to see all serial output that comes after the initial boot, useful for monitoring the OS init flow.

- Go to “Advanced” tab → “Serial Port Console Redirection” → Under “Serial Communication via IPMI COM”.
- Set “Console Redirection” to “Enabled”.




At this point, make sure to disconnect the VGA monitor and USB keyboard, or else the following error may appear:

TSC_DEADLINE disabled due to Errata; Please update microcode to version : 0xffffffff or later

3. Log in as admin and use admin as password, using IPMI.

```
ipmitool -I lanplus -H <IP Address> -U admin -P admin sol activate
```

4. Connect the management Ethernet cable to LAN3 (second port from the left) on the back panel of the appliance.



5. Go through the Gateway Management configuration wizard.
IP Configuration by DHCP

Wizard Session Display (Example)	Comments
Do you want to use the wizard for initial configuration? yes	This configuration must be performed the first time the gateway is operated or after resetting the gateway to the factory defaults. Type "y" and then press <Enter>.
Step 1: Hostname? [gateway-1]	To accept the default hostname, press <Enter>. Otherwise, type a different hostname and press <Enter>.
Step 2: Use DHCP on mgmt0 interface? [yes]	Perform this step to obtain an IP address for the gateway (mgmt0 is the management port of the gateway). • Typing "yes" will have the DHCP server assign the IP address • Typing "no" (no DHCP) will offer the use of the "zeroconf" configuration or not. For the use of Zeroconf, type "yes" and the session will continue. If "no" (no Zeroconf) is typed, enter a static IP and the session will continue.
Step 3: Enable IPv6 [yes]	Perform this step to enable IPv6 on management ports. • Type "yes" to enable enable IPv6. • Type "no" to not enable IPv6 (Step 4 will be skipped)
Step 4: Enable IPv6 autoconfig (SLAAC) on mgmt0 interface	Perform this step to enable Stateless address autoconfig on external management port. • Type "yes" to enable • Type "no" to disable
Step 5: Use DHCPv6 on mgmt0 interface? [yes]	Perform this step to enable DHCPv6 on the MGMT0 interface.
Step 6: Enable password hardening?	Perform this step to enable/disable password hardening on your machine. If enabled, new passwords will be checked upon configured restrictions. If you wish to enable it, type "yes" and press . If you wish to disable it, enter "no"
Step 7: Admin password (Must be typed)? <new_password>	To avoid illegal access to the machine, type a password and press <Enter>. An admin password must be entered upon initial configuration. Due to California Senate Bill No. 327, this stage is required and cannot be skipped.
Step 8: Confirm admin password? <new_password>	Confirm the password by re-entering it. Note that password characters are not printed.

Wizard Session Display (Example)	Comments
Step 9: Monitor password (Must be typed)? <new_password>	To avoid illegal access to the machine, please type a password and then press <Enter>. An admin password must be entered upon initial configuration. Due to California Senate Bill No. 327, this stage is required and cannot be skipped.
Step 10: Confirm monitor password? <new_password>	Confirm the password by re-entering it. Note that password characters are not printed.
You have entered the following information: Hostname: <gateway name> Use DHCP on mgmt0 interface: yes Enable IPv6: yes Enable IPv6 autoconfig (SLAAC) on mgmt0 interface: yes Enable DHCPv6 on mgmt0 interface: no Enable password hardening: yes Admin password (Enter to leave unchanged): (CHANGED) To change an answer, enter the step number to return to. Otherwise hit <enter> to save changes and exit. Choice: <Enter> Configuration changes saved. To return to the wizard from the CLI, enter the "configuration jump-start" command from configuration mode. Launching CLI... <gateway name> [standalone: master] >	The wizard displays a summary of choices and then asks to confirm the choices or to re-edit them. - Press <Enter>, to save changes and exit - Enter the relevant configuration step number, to edit any of the choices To run the command "configuration jump-start", Config mode must be used.

- Check the mgmt0 interface configuration before attempting a remote connection (e.g., SSH) to the gateway. Specifically, verify the existence of an IP address.

```

gateway # show interfaces mgmt0

Interface mgmt0 status:
  Comment      :
  Admin up     : yes
  Link up      : yes
  DHCP running : yes
  IP address   : 10.7.148.61
  Netmask      : 255.255.0.0
  IPv6 enabled : yes
  Autoconf enabled: no
  Autoconf route : yes
  Autoconf privacy: no
  DHCPv6 running : no
  IPv6 addresses : 1

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

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

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

Tx:
  5139846 bytes
  28452 packets

```

```

0 discards
0 errors
0 overruns
0 carrier
0 collisions
1000 queue len

```

4.1.1 Rerunning the Wizard

To rerun the wizard, do the following:

1. Enter config mode.

```

gateway > enable
gateway # config terminal

```

2. Rerun the wizard.

```

gateway (config) # configuration jump-start

```

4.2 Starting the Command Line Interface (CLI)

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

```

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

```

3. Log in to the gateway (default username and password are both “admin”).
4. Read and accept the EULA, when prompted.
5. Once the following prompt appears, the system is ready to use.

```

a. NVIDIA Gateway

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

gateway >

```



If firmware was upgraded, firmware boot bar will appear and the CLI will be blocked until firmware upgrade is complete.



The CLI will be blocked until InfiniBand virtual interfaces are created. The following message will appear: “[Creating VFs](#)”.

```

csl-skyway login: admin
Password:
Last login: Sun Mar  7 16:05:03 IST 2026 from 10.228.129.138 on pts/0
Number of total successful connections since last 1 days: 7
Your capability and password have been changed

Mellanox Gateway

System is initializing!
This may take a few minutes

creating VFS: | | | |

```



5 Networkwide Deployment Guidelines

Skyway GA100 is an appliance-based InfiniBand-to-Ethernet gateway, enabling Ethernet-based communications to access the InfiniBand datacenter and vice versa. The following section describes networkwide guidelines and provides a specific example when using a NVIDIA Ethernet switch running NVIDIA Onyx™ operating system.



Ensure the Subnet Manager is running in the InfiniBand cluster with IPoIB enabled.



Make sure that validation of network connectivity is be done strictly between Ethernet and InfiniBand host. Ping with test IPs belonging to the Skyway are mostly invalid and should not be used for validation.




The IPs and subnet masks described in the diagram below are provided as an example. Actual IP and subnet masks may differ in actual customer deployments.



The Skyway solution mandates a dedicated subnet between the appliance and the layer-3 Ethernet switch/router. This subnet is not visible from outside and can be configured to be ANY class-C subnet, including IPs which are not part of the customer's IP pool.

Skyway appliance MUST be directly attached to L3 Ethernet switch(es)/router(es) (at least one) to complement the function of an IPoIB router. The role of divisions is as follows: Skyway forwards IPoIB traffic to a single IP address/MAC on the Ethernet side; the directly-attached L3 Ethernet switch(es)/router(es) provide Ethernet-side load balancing functionality (such as MLAG) and routing.

All Skyway devices with a single HA domain share the same virtual IP.



5.1 General Networkwide Guidelines

5.1.1 Ethernet Guidelines

The connection between the Skyway and the Ethernet router requires configuring a LAG with active LACP on the Ethernet router (see step 3 in the "Configuring IP Addresses and Routes" below).

For increased resiliency, it is recommended to configure Ethernet routers in an [MLAG configuration](#).



Make sure MTU on the Ethernet router connected to Skyway is at least 2 bytes smaller than the InfiniBand IPoIB MTU configuration (e.g., set InfiniBand MTU to 4092 and Ethernet MTU to 4090).

5.1.2 InfiniBand Guidelines



All InfiniBand ports must be connected to the same InfiniBand fabric.

5.1.2.1 Subnet Manager Configuration



Ensure the Subnet Manager is running in the InfiniBand cluster with IPoIB enabled.

Virtualization must be enabled by the Subnet Manager (SM). It is recommended to unlimit the maximum number of ports that are processed simultaneously.

If opensm runs on an InfiniBand switch, configure the following:

```
switch (config) # ib sm virt enable
switch (config) # ib sm virt-max-ports-in-process 0
```

If opensm runs on a host, add the following lines to the opensm.conf (by default at /etc/opensm/opensm.conf):

```
# Virtualization support
# 0: Ignore Virtualization - No virtualization support
# 1: Disable Virtualization - Disable virtualization on all
# Virtualization supporting ports
# 2: Enable Virtualization - Enable (virtualization on all
# Virtualization supporting ports
virt_enabled 2

# Maximum number of ports to be processed simultaneously
# by Virtualization Manager (0 - process all pending ports)
virt_max_ports_in_process 0
```

5.2 Configuring IP Addresses and Routes



The first port (port #1) of each HCA is an InfiniBand port and the second port (port #2) is an Ethernet port; therefore, the configuration of an InfiniBand "device/port" value should be "x/1" and the configuration of an Ethernet "device/port" value should be "x/2".

For example, for HCA #7, the configuration of the InfiniBand port is 7/1 and of the Ethernet port is 7/2.



IP addresses, subnet masks, port numbers, and interface names are used as an example and may vary according to the actual connectivity of the customer's deployment.

1. On the relevant InfiniBand nodes, configure an IP address on each InfiniBand port designated for the Skyway deployment (e.g., ib0). In addition, configure a default route with the Skyway IP as next hop (in this example, 1.1.1.3).

```
# ifconfig ib0 1.1.1.2/24
# ip route add 0/0 via 1.1.1.3
```



The NVIDIA Skyway IP which is configured as next hop should match the virtual IP of the InfiniBand port channel of the NVIDIA Skyway appliance.

2. Access and configure an IP address on the gateway's Ethernet and InfiniBand ports and configure a virtual IP address on the InfiniBand port. In addition, configure a route to the customer's Ethernet networks via the IP assigned on the Ethernet router's port (in this example, 2.2.2.1).

```
gateway > enable
gateway # configure terminal
gateway (config) # interface ib port-channel 1 ip address 1.1.1.1/24
gateway (config) # interface ib port-channel 1 virtual ip address 1.1.1.3/24
gateway (config) # interface ethernet port-channel 1 ip address 2.2.2.2/24
gateway (config) # ip route 0.0.0.0/0 2.2.2.1
```

3. Detect the ports on the Ethernet router that are connected to the gateway, assign the LAG to a VLAN, and configure the IP address on the VLAN interface. In addition, configure a route to the IPoIB network via the IP assigned to the gateway's Ethernet port-channel (in this example, 2.2.2.2).


Below is an example using a NVIDIA-Onyx-based switch with port 1/1 connected to the gateway.

In this example, ports 1-8 on the router (see line 7) are connected to the 8 Ethernet ports on Skyway.

```
eth_router > enable
eth_router # configure terminal
eth_router (config) # ip routing
eth_router (config) # lacp
eth_router (config) # interface port-channel 1
eth_router (config) interface port-channel 1) # exit
eth_router (config) # interface ethernet 1/1-1/8 channel-group 1 mode active
eth_router (config) # vlan 2
eth_router (config) vlan 2) # exit
eth_router (config) # interface port-channel 1 switchport access vlan 2
eth_router (config) # interface vlan 2 ip address 2.2.2.1 /24
eth_router (config) # ip route 1.1.1.0 /24 2.2.2.2
```



Note that the above connection describes a connection between Skyway and a single Ethernet router with LAG. It is possible to also connect to two Ethernet routers in an MLAG configuration. For more information, see the following community post for MLAG configuration on NVIDIA Onyx-based switches: <https://enterprise-support.nvidia.com/s/article/how-to-configure-mlag-on-the-switches>.



Additional commands might be needed to ensure the MTU restrictions described in [Ethernet Guidelines](#) above.

4. To ensure proper deployment, ping between a host in the subnet 192.168.1.0/24 and the InfiniBand host with IP 1.1.1.2 should be successful.



Make sure that validation of network connectivity is done strictly between Ethernet and InfiniBand host.
Ping with test IPs belonging to the Skyway are mostly invalid and should not be used for validation.

List of **valid** ping conditions:

- Ping from InfiniBand server to Skyway InfiniBand VIP
- Ping from InfiniBand server to Skyway Eth port channel interface IP
- Ping from InfiniBand server to Eth router interface IP
- Ping from Eth server to Skyway Eth port channel interface IP
- Ping from Eth server to Eth router interface IP

List of **invalid** ping conditions:

- Ping from InfiniBand server to Skyway InfiniBand port channel interface IP will fail or be intermittent
- Ping from Ethernet server to Skyway VIP should fail or be intermittent
- Ping from Ethernet server to Skyway InfiniBand port channel interface IP should fail or be intermittent
- Ping from Skyway Ethernet port-channel are unreliable: ping can fail (response sent to the other Skyway) or succeed (response sent to the same Skyway)

5.3 Configuring High Availability (HA)

This section explains how to configure a HA cluster with multiple appliances.

5.3.1 Before Configuring HA



- For all appliances in the HA cluster, the MLNX-GW version must be the same.
- For all appliances in the HA cluster, the Ethernet management interfaces must be in the same L2 subnet.
- The Skyway appliances configured in HA mode must be connected to either Ethernet L3-dedicated switch or Ethernet L2 where all ports connected to Skyway are configured as router ports.

- Before configuring HA, each appliance should be configured according to the "Configuring IP Addresses and Routes" section above.
- Virtual IP configuration and Ethernet port channel configuration must be identical for all appliances in the HA cluster.

Example of configuration that needs to be identical for all appliances:

Skyway A:

```
gateway(config) # interface ib port-channel 1 virtual ip address 1.1.1.3/24
gateway(config) # interface ethernet port-channel 1 ip address 2.2.2.2/24
```

Skyway B:

```
gateway(config) # interface ib port-channel 1 virtual ip address 1.1.1.3/24
gateway(config) # interface ethernet port-channel 1 ip address 2.2.2.2/24
```

- The ib port channel IP address may be different between the appliances in the HA cluster:

Skyway A:

```
gateway(config) # interface ib port-channel 1 ip address 1.1.1.1/24
```

Skyway B:

```
gateway(config) # interface ib port-channel 1 ip address 1.1.1.4/24
```

- Make sure that all Ethernet interfaces that are connected to Skyway appliances in the same HA cluster are connected through an Ethernet MLAG or LAG configuration. Below is an example of MLAG and MAGP configuration on Ethernet switches connected to Skyway appliances.

```
eth_router > enable
eth_router # configure terminal
eth_router (config) # protocol mlag
eth_router (config) # lacp
eth_router (config) # vlan 999
eth_router (config) # exit
eth_router (config) # interface vlan 999 ip address 192.17.10.3/24 primary
eth_router (config) # interface port-channel 1 # exit
eth_router (config) # interface ethernet 1/1-1/4 channel-group 1 mode active
eth_router (config) # interface port-channel 1 ipl 1
eth_router (config) # interface vlan 999 ipl 1 peer-address 192.17.10.2
eth_router (config) # mlag-vip GW-HA ip 10.10.252.10 /16 force
eth_router (config) # no mlag shutdown
eth_router (config) # interface mlag-port-channel 101
eth_router (config) # interface mlag-port-channel 101 # exit
eth_router (config) # interface ethernet 1/19-1/26 mlag-channel-group 101 mode active
eth_router (config) # interface mlag-port-channel 101 no shutdown

eth_router (config) # ip routing
eth_router (config) # vlan 101
```

```

eth_router (config) vlan 101) # exit
eth_router (config) # interface vlan 101 ip address 2.2.2.252/24 primary
eth_router (config) # interface mlag-port-channel 101 switchport access vlan 101
eth_router (config) # protocol mlagp
eth_router (config) # interface vlan 101 mlagp 101
eth_router (config) interface vlan 101 mlagp 101) # ip virtual-router address 2.2.2.254
eth_router (config) interface vlan 101 mlagp 101) # ip virtual-router mac-address AA:BB:CC:00:01:01
eth_router (config) # ip route vrf default 172.0.0.0/8 2.2.2.2

```

Below is an example of LAG configuration on Ethernet switch connected to Skyway appliances. Ports 1–8 on the router are connected to the 8 Ethernet ports on the first Skyway appliance and ports 11–18 on the router are connected to the 8 Ethernet ports on the second Skyway appliance.

```

eth_router > enable
eth_router # configure terminal
eth_router (config) # ip routing
eth_router (config) # lacp
eth_router (config) # interface port-channel 1
eth_router (config) interface port-channel 1) # exit
eth_router (config) # interface ethernet 1/1-1/8 channel-group 1 mode active
eth_router (config) # interface ethernet 1/11-1/18 channel-group 1 mode active
eth_router (config) # vlan 2
eth_router (config) vlan 2) # exit
eth_router (config) # interface port-channel 1 switchport access vlan 2
eth_router (config) # interface vlan 2 ip address 2.2.2.1 /24
eth_router (config) # ip route 1.1.1.0 /24 2.2.2.2

```



Even if working on a single Skyway appliance system, it is recommended to configure the appliance to have High Availability configuration on the system. This will allow to easily scale the topology in the future without needing to change a single Skyway appliance configuration. See section "Configuring HA on Skyway Appliance" below for configuration details.

5.3.2 Configuring HA on Skyway Appliance

1. Configure HA on the gateway. Configure HA on each Skyway appliance that is going to be a part of the HA cluster.

All Skyway appliances must share the same HA domain.

Skyway A:

```

gateway (config) # gw ha 1
Warning! Configuration is about to be saved and the system will be reloaded.
Type 'YES' to confirm the HA domain id change: YES

```

Skyway B:

```

gateway (config) # gw ha 1
Warning! Configuration is about to be saved and the system will be reloaded.
Type 'YES' to confirm the HA domain id change: YES

```



After this step, the Skyway appliances will be rebooted.

2. Once all systems complete the initialization, verify that all Skyway appliances were added properly to the HA cluster by running "show gw ha" from one of the Skyway appliances. Verify domain ID appears as configured and all Skyway appliances appear in the output of the command.

```

gateway (config) # show gw ha

Global HA state:

```

```

GW domain ID : 3
Active HA nodes: 3
Master name : skyway-7


HA domain nodes information:
Name : skyway-8
GW Operational state: active
System guid : b8ce:f603:0075:6eda
Priority : 100

Name : skyway-64
GW Operational state: active
System guid : b8ce:f603:0068:7e8a
Priority : 100


Name : skyway-7 <--- (local node)
GW Operational state: active
System guid : b8ce:f603:0075:6efa
Priority : 100

```

5.3.2.1 High Availability LAG/MLAG Setup



5.3.2.2 Skyway Connectivity to the Ethernet Using L2 Ethernet Switches



5.4 Configuring Partition Keys (PKEYs)

1. Configure partition keys on the Skyway appliance.

```
#### pkey_id = [1, 7ffe] (hex)
```

2. Configuring partition keys on the subnet manager.

```
ib switch (config) # ib partition <partition name> pkey 0x<pkey_id>
ib switch (config) # ib partition <partition name> ipoib
ib switch (config) # ib partition <partition name> member ALL type full
```

Example:

```
ib switch (config) # ib partition pkey_0x1 pkey 0x1
ib switch (config) # ib partition pkey_0x1 ipoib
ib switch (config) # ib partition pkey_0x1 member ALL type full
```

3. Configure PKEY on the InfiniBand host.

```
#### pkey_full = pkey_id (hex) + 8000 (hex)
echo "0x<pkey_full>" > /sys/class/net/ib<interface_number>/create_child
ifconfig ib<interface_number>.<pkey_full> up
ifconfig ib<interface_number>.<pkey_full> <Ib host pkey port ip> netmask <netmask>
ip route add <eth subnet> via <Gw pkey port virtual ip>
ip route add <gw eth subnet> via <Gw pkey port virtual ip>
```

Example:

```
echo "0x8001" > /sys/class/net/ib1/create_child
```

```

ifconfig ibl.8001 up
ifconfig ibl.8001 111.222.62.11 netmask 255.255.255.0
ip route add 2.2.2.0/24 via 111.222.62.1
ip route add 192.168.1.0/24 via 111.222.62.1

```



Note that this example shows how to configure PKEY that has an ID of 0x1 (0x8001 in hex == 8000+pkeyID). When configuring other PKEYs on the InfiniBand host, make sure to add 8000 to the PKEY ID.

4. Configure the Skyway appliance.

```

gateway (config) # interface ib port-channel 1 pkey 0x<pkey_id>
gateway (config) interface ib port-channel 1 pkey 0x<pkey_id>) # exit
gateway (config) # interface ib port-channel 1 pkey 0x<pkey_id> ip address <Gw pkey port ip> <netmask>
gateway (config) # interface ib port-channel 1 pkey 0x<pkey_id> virtual ip address <Gw pkey port virtual
ip> <netmask>
gateway (config) # configuration write

```

Example:

```

gateway (config) # interface ib port-channel 1 pkey 0x1
gateway (config) interface ib port-channel 1 pkey 0x1) # exit
gateway (config) # interface ib port-channel 1 pkey 0x1 ip address 111.222.62.2 255.255.255.0
gateway (config) # interface ib port-channel 1 pkey 0x1 virtual ip address 111.222.62.1 255.255.255.0
gateway (config) # configuration write

```

5. Add routes on the Ethernet host.

```

ip route add <pkey subnet>/<netmask> via <eth router port ip>

```

Example:

```

ip route add 111.222.62.0/24 via <eth router port ip>

```

6. Add routes on the Ethernet router.

```

ip route <pkey subnet>/<netmask> <gw eth pc ip>

```

Example:

```

ip route 111.222.62.0 /24 <gw eth pc>

```

6 System Monitoring

6.1 Front Panel Monitoring Components




6.1.1 Power-On LED

There is one I/O LED (green) on the front panel to indicate if the system is powered.

LED State	Color	Description
On	Green	System is turned on
Blinking	Green	System is under S4 state
Off	N/A	Power off

6.2


Rear Panel LEDs



6.2.1 LAN Interfaces LEDs

6.2.1.1 LAN3/LAN4 Rear I/O LED Interface

There are two I/O LEDs (green and amber) to indicate LAN link and activity.




Left LED	Right LED	Description
Off	Green	10M bps linked
Off	Blinking Green	10M bps active
Amber	Green	100M bps linked
Amber	Blinking Green	100M bps active

Left LED	Right LED	Description
Green	Green	1 G bps linked
Green	Blinking Green	1 G bps active
Off	Off	No link

6.2.1.2 LAN1/LAN 2 Rear I/O LED Interface

There are two I/O LEDs (green and amber) to indicate LAN link and activity.



Left LED	Right LED	Description
Amber	Green	1 G bps linked
Amber	Blinking Green	1 G bps active
Green	Green	10G bps linked
Green	Blinking Green	10G bps active
Off	Off	No link

6.2.2 Power Module LED

There are two I/O LEDs (amber and green) to indicate the power module state.


LED State	Color	Description
Blinking	Amber	Power supply warning event
Blinking	Green	AC present standby output on
On	Amber	AC unplug to this module or power supply critical event
On	Green	Power supply DC output ON and OK
Off	Off	No AC power to both power modules

6.2.3 Network Interface Cards LEDs

There are two I/O LEDs per port:

- LED 1 and 2: Bi-color I/O LED which indicates link status. LED behavior is described below for Ethernet and InfiniBand port configurations.

- LED 3 and 4: Reserved for future use.



LED1 and LED2 Link Status Indications (Physical and Logical)—Ethernet Protocol

LED Color and State	Description
Off	A link has not been established.
Blinking amber	1 Hz blinking amber occurs when a beacon command for locating the adapter card running . 4 Hz blinking amber indicates a problem with the physical link.
Solid green	Indicates a valid link with no active traffic.
Blinking green	Indicates a valid logical link with active traffic.

LED1 and LED2 Link Status Indications (Physical and Logical)—InfiniBand Protocol


LED Color and State	Description
Off	A physical link has not been established.
Solid amber	Indicates an active physical link.
Blinking amber	1 Hz blinking amber occurs when a beacon command for locating the adapter card is running. 4 Hz blinking amber indicates a problem with the physical link.
Solid green	Indicates a valid logical (data activity) link with no active traffic.
Blinking green	Indicates a valid logical link with active traffic.

7 System Maintenance

This chapter contains the installations and un-installation instructions of the following customer replaceable units.

7.1 Power Supply Units

Skyway-3 is equipped with two replaceable power supply units (PSU) that work in a redundant configuration. The figure below shows the power side of the system which includes a hot-swap PSU.



Item	Description
1	Power Socket
2	Extraction handle
3	PSU Status LEDs

7.1.1 Extracting and Inserting the Power Supply Unit



The power supply is only hot-swappable if you have a redundant system with two power supplies installed. If there is only one power supply installed, before removing or replacing the power supply, first take the appliance out of service, turn off all peripheral devices connected to the system, turn off the system by pressing the power button, and unplug the AC power cord from the system or wall outlet.

The PSU can be replaced in case it fails.



To extract the PSU, do the following:

Step 1. If a filler panel is installed, remove the filler panel.

Step 2. If a PSU is installed, grab the handle with your thumb pointing toward the latch. Push the latch with your thumb towards the handle while you pull the PSU out of the appliance.



To insert a PSU, do the following:

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



Do not run the appliance with openings due to missing parts. This may cause overheating due to improper airflow.

Step 2. Insert the PSU by sliding it into the opening, until a slight resistance is felt.

Step 3. Continue pressing the PSU until it seats completely. The latch will snap into place, confirming proper installation.

Step 4. Insert the power cord into the supply connector.

Step 5. Insert the other end of the power cord into an outlet of the correct voltage.

7.2 Slide Rail Kit



To disassemble the appliance from the rack, do the following:

Step 1. Shut down the appliance.

Step 2. Unplug and remove all connectors.

Step 3. Unplug all power cords.

Step 4. Remove the ground wire.

Step 5. Unscrew the 2 center bolts from inside the handles.

Step 6. Slide the appliance from the rack.

Step 7. Remove the rail slides from the rack.



To install the slide rail, do the following:

- Please refer to [Slide Rail Installation](#) section.

8 Troubleshooting

8.1 General Troubleshooting

Issue	Resolution
System Status LED is RED	Unplug the appliance and call your NVIDIA Networking representative.
Power Supply Unit Status LED is not lit or is RED	1. Check that the power cable is plugged into a working outlet. 2. Check that the power cable has a voltage within the range of 100–240 volts AC. 3. Remove and reinstall the power cable. 4. Remove and reinstall the PSU.
The Power Button w/ Integrated LED for the appliance shuts off	1. Check that there is adequate ventilation. 2. Make sure that there is nothing blocking the front or rear of the chassis and that the fan modules and ventilation holes are not blocked (especially look for dust over the holes). 3. If you find dust blocking the ventilation holes, it is recommended to clean the fan unit and remove the dust from the front and rear panels of the appliance using a vacuum cleaner.
The Activity LEDs do not come on	Check if the NVIDIA Skyway-3 appliance has been started.
The appliance is off	1. Press the Power Button w/Integrated LED If that does not work, do the following: 1. Unplug the appliance. 2. Wait 5 minutes. 3. Plug in the appliance, and press the Power Button w/Integrated LED. 4. If the appliance does not come on, check the power supplies. 5. If the appliance comes on, use the NVIDIA Skyway-3 management software to determine the cause of the shutdown. 6. Check the temperature. 7. Check the fan status.

9 Technical Specifications

9.1 MGA400-XS2 Specifications



Physical	Dimensions (W x H x D): 438.4 x 88 x 900 (17.25" x 3.46" x 35.43")		
	Weight: 39.2kg		
	Mounting: MGX		
Protocol Support	InfiniBand: IBTA v1.3 ^a Auto-Negotiation: SDR (2.5Gb/s per lane), DDR (5Gb/s per lane), EDR (25Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane), HDR (50Gb/s per lane) port		
	Ethernet: 200GBASE-CR4, 200GBASE-KR4, 200GBASE-SR4, 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 50GBASE-R2, 50GBASE-R4, 25GBASE-R		
	Data Rate	InfiniBand	SDR/EDR/HDR100/HDR/NDR
		Ethernet	25/50/100/200/400 Gb/s
	Gen3: SERDES @ 8.0GT/s, x16 lanes (2.0 and 1.1 compatible)		
Power	Voltage: Monitors for CPU Cores, +3.3 V, +5 V, +12 V, +5 V Standby, VBAT		
	80 PLUS Platinum 2+1 redundant power supply 1000 W @ 100 ~ 127 V 2000 W @ 200 ~ 240 V		
	Detailed Power Consumption Per Interface		
	Interface	Power Consumption Per Interface	Total Power Consumption Per Interface
	2x CPU	2x185W	370W
	16x DDR5 32G	16x10W	160W
	1x SSD NVME 1.92TB	1x25W	25W
	6x fans	6x73W	438W
	8x ConnectX-8 Network Cards	8x60W	480W
	1x BMC	1x2W	2W
	2x USB	2.5W	5W
	Other server components		3W
	Maximum Power Consumption of the entire server		~1652W
Environmental	Temperature	Operating	10 to 35° C
		Non-operating	-30 to 60° C
	Humidity	Operating relative humidity	8~80% @ 40° C (non-condensing)
		Non-operating relative humidity	8~90% @ 40° C (non-condensing)
	Vibration (5 ~ 500 Hz)	0.21 GRMs	

	Shock	5G (with 10ms duration, half sine wave)
Thermal	Airflow	30 CFM
	Maximum heat dissipation	5153 BTU/hour

9.2 System Dimensions



All dimensions are in millimeters. All the mechanical tolerances are +/- 0.1mm.



9.3 Thermal Threshold Definitions

There are two thermal threshold definitions for NVIDIA Skyway-3 which impact the overall system operation state:

1. Critical—When the device crosses this temperature, the firmware will automatically shut down the device. This temperature threshold is set from the BIOS (Advanced > IT8528 HW

Monitor > CPU ACPI Shutdown Temperature). The temperature threshold can be configured from 50 to 110°C.

2. Emergency—The temperature threshold is set by the CPU's internal thermal trip. It is impossible to change the temperature value through a software interface.

10 Inventory Information

The system's inventory parameters (such as the serial number or part number) can be extracted from labels on the system's bottom side.



型号/型號 **Model: MGA100** 产品名称: 交换机 / 產品名稱: 交換器

Electrical Rating: 1x/2x 100-127Vac, 47/63Hz, 12A

輸入規格 / 輸入規格 **200-240Vac, 47/63Hz, 10A**





CAN ICES-3 (A)NMB-3(A)

Pera Argentina AC
200-240V, 47/63Hz

警告使用者：這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾。在這種情況下，使用者會被要求採取某些適當的對策。

此為A級產品，在生活環境中，該產品可能會造成無線電干擾。在這種情況下，可能需要用戶對其干擾採取切實可行的措施。

European Regulatory Address:
Mellanox Technologies Denmark ApS,
Lødrborg Allé 136B,
4000 Tørløse, Denmark

Patented: www.mellanox.com/page/patents

This device complies with part 15 of the FCC rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Laite on tillettävä suojakoskettimilla varustettuun pistorasialaan. Apparat må tilkoples jordet stikkontakt. Apparatet skall anslutas till jordat uttag.

CAUTION Risk of electric shock and energy hazard. Disconnecting one power supply disconnects only one power supply module. To isolate the unit completely disconnect all power supplies.	ATTENTION Risque de choc et de danger électriques. Le débranchement d'une seule alimentation stabilisée ne débranche uniquement qu'un module "Alimentation Stabilisée". Pour isoler complètement le module on cause, il faut débrancher toutes les alimentations stabilisées.	注意安全 注意：此产品可能为两个电源供电，单电源可以正常工作。在维修的时候一定要关闭所有电源	ACHTUNG Gefahr des elektrischen Schocks. Entfernen des Netzsteckers eines Netz-teils spannungsfrei. Um alle Einheiten spannungs-frei zu machen sind die Netzstecker aller Netzteile zu entfernen.
--	---	--	---

DO NOT REMOVE THE COVER www.mellanox.com

S/N: MT2004X10163



P/N: MGA100-HS2

Rev: A1

Made in Israel

2020-06-23



11 Field Replaceable Units

Ordering Number	Part Description
MGA400-PS	Power supply for NVIDIA Skyway Infiniband-to-Ethernet appliance
MGA400-RKIT	Rail kit for NVIDIA Skyway Infiniband-to-Ethernet appliance
SSD FRU	

12 Revision History

Date	Description of Changes
July. 2026	First 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 its affiliates 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

© 2026 NVIDIA Corporation & affiliates. All Rights Reserved.

