

NVIDIA Skyway InfiniBand-to-Ethernet Gateway User Manual

Table of Contents

| 1 | Introduction | 5 |
|-------|--|----|
| 1.1 | Product Overview | 5 |
| 1.1.1 | NVIDIA Skyway Highlights | 5 |
| 1.2 | Main System Components | 6 |
| 1.2.1 | Network Interface Cards | 6 |
| 1.2.2 | Power Supply Units | 6 |
| 1.2.3 | Fans | 6 |
| 1.3 | Package Contents | 7 |
| 1.3.1 | Rail Kit Package Contents | 7 |
| 1.4 | System Features | 8 |
| 1.4.1 | InfiniBand-to-Ethernet Gateway Operational Description | 8 |
| 1.4.2 | Load Balancing and High Availability Operational Description | 8 |
| 1.5 | Operating System | 9 |
| 1.6 | Certifications | 9 |
| 2 | System Layout and Interfaces | 10 |
| 2.1 | NVIDIA Skyway Front and Rear Panel | 10 |
| 2.2 | Interfaces Detailed Description | 11 |
| 2.2.1 | Power-On LED | 11 |
| 2.2.2 | USB 3.0 Interfaces | 11 |
| 2.2.3 | PCIe Gen 3.0 Slots | 11 |
| 2.2.4 | LAN Interfaces | 11 |
| 2.2.5 | Redundant Power Module | 12 |
| 2.2.6 | Fans Modules | 12 |
| 3 | Hardware Installation | 13 |
| 3.1 | NVIDIA Skyway Installation | 13 |
| 3.1.1 | Safety Warnings | 13 |
| 3.1.2 | System Requirements | 14 |
| 3.1.3 | Unpacking the Package | 14 |
| 3.1.4 | Rack Mounting | 15 |
| 3.2 | Cable Installation | 19 |
| 3.2.1 | Power Cable | 19 |
| 3.2.2 | ConnectX-6 Networking Cards Cables | 19 |

| 3.3 | System Connectivity 21 |
|-------|--|
| 3.4 | Initial Power-On |
| 4 | Configuring the Gateway for the First Time 23 |
| 4.1 | Gateway Initialization23 |
| 4.1.1 | Rerunning the Wizard27 |
| 4.2 | Starting the Command Line Interface (CLI) |
| 5 | Networkwide Deployment Guidelines |
| 5.1 | Configuring High Availability (HA) 29 |
| 5.1.1 | Before Configuring HA29 |
| 5.1.2 | Configuring HA on Skyway Appliance30 |
| 5.2 | Configuring Partition Keys (PKEYs) |
| 6 | System Monitoring |
| 6.1 | Front Panel Monitoring Components 35 |
| 6.1.1 | Power-On LED |
| 6.2 | Rear Panel LEDs |
| 6.2.1 | LAN Interfaces LEDs36 |
| 6.2.2 | Power Module LED36 |
| 6.2.3 | Network Interface Cards LEDs37 |
| 7 | System Maintenance |
| 7.1 | Power Supply Units |
| 7.1.1 | Extracting and Inserting the Power Supply Unit |
| 7.2 | Slide Rail Kit |
| 8 | Troubleshooting |
| 8.1 | General Troubleshooting 41 |
| 9 | Technical Specifications |
| 9.1 | MGA100-HS2 Specifications |
| 9.2 | System Dimensions |
| 9.3 | Thermal Threshold Definitions 43 |
| 10 | Inventory Information |
| 11 | Field Replaceable Units |
| 12 | Revision History |

 ${\rm (i)}$ You can download a PDF version of the full guide here.

About this Manual

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

Ordering Part Numbers

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

| NVIDIA SKU | Legacy OPN | Marketing Description |
|--------------------|------------|--|
| 920-9B020-00FA-0D2 | MGA100-HS2 | NVIDIA Skyway InfiniBand to Ethernet Gateway Appliance, 8x IB and 8x Ethernet ports, 2U server, 8x ConnectX-6 VPI dual-port HDR/200GbE adapters, 2 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 InfiniBand-to-Ethernet gateway. This document contains the complete product overview, installation and initialization instructions, and product specifications.

1.1 Product Overview

NVIDIA Skyway 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.



1.1.1 NVIDIA Skyway Highlights

| Componenet | MGA100-HS2 |
|---------------------------|---|
| Form Factor | 2U rackmount: 19" |
| Weight | NVIDIA Skyway gateway: 21kg The gateway with ACC and package: 32kg |
| PCIe Cards | 8x NVIDIA® ConnectX®-6 VPI 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 100Gb/s bi-directional per port |
| Power Supplies | 2x AC power supplies |

⁽i) This document is preliminary and subject to change.

1.2 Main System Components

The NVIDIA Skyway 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

NVIDIA Skyway 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 is equipped with two 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.

A 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 is equipped with one fan per power supply unit on the rear panel of the appliance.

1.2.3.2 Internal Fans

NVIDIA Skyway 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 2U system | |
| Slide Rail Kit | 1 | 1U/2U 36" slide kit pair for NVIDIA Skyway | |
| Power Cables 2 250V 10A 1830MM C14 TO C13 power cable | | 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 |
|--------------------------|-----------------|--|
| | СРИ Туре | Dual Intel LGA3647 Xeon Scalable processor (up to 140W TDP) |
| Memory Type | | Supports DDR4 2133/2400/2666 MHz ECC-REG Modules |
| | Memory Size | 4 x 16GB DDR4 2666MHz |
| | Memory Voltage | 1.2V |
| | Error Detection | Corrects single-bit errors Detects double-bit errors (using ECC memory) |
| Rear I/O Panel | USB | 4x USB 3.0 |

| | RJ-45/LAN | 4 x RJ-45 LAN ports: • 2x 10GbE • 2x 1GbE/IPMI-LAN | |
|------------------|-------------------------------|--|--|
| On-board Devices | EC | TE 8528E chip provides motherboard, RS-232, and hardware monitor functions | |
| | BMC | Sharing with the LAN 1/4. | |
| Expansion slots | PCI-Express | 8x network interface cards | |
| Cooling | Chassis Fan | 2x 4-pin 80x38 high-speed fans for CPU 4x 4-pin 80x38 high-speed fans for expansion cards | |
| | PSU Fans | One fan per power supply unit | |
| PC Health | Voltage | Monitors for CPU Cores, +3.3 V, +5V, +12V, +5V standby, VBAT | |
| Monitoring | 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 <u>www.nvidia.com/en-us/</u> <u>networking</u>. In the main menu, click on PRODUCTS \rightarrow INFINIBAND \rightarrow GATEWAY & ROUTERS SYSTEMS \rightarrow select the desired product page.

1.4.1 InfiniBand-to-Ethernet Gateway Operational Description

NVIDIA Skyway 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 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 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 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 or100Gb/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 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 <u>NVIDIA MLNX-GW User Manual for NVIDIA Skyway</u> 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 Netowrking website at <u>http://www.mellanox.com/page/</u><u>environmental_compliance</u>.

2 System Layout and Interfaces

📜 Unknown macro: 'easy-heading-free'

The figures below show the front and rear sides of NVIDIA Skyway. 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 Front and Rear Panel



| Item | Interface | Description |
|------|------------------------|--|
| 1 | Power-On LEDs | Bezel LED-lights when the system is powered on |
| 2 | COM port | Serial DB9 to RJ45 RS232 port-not active |
| 3 | USB Interfaces | 4 x USB 3.0 ports at the rear panel |
| 4 | PCIe Gen 3.0/4.0 Slots | 8x ConnectX-6 VPI adapter cards |
| 5 | LAN1 | RJ-45 LAN port-10G IPMI-LAN2 |
| 6 | LAN2 | RJ-45 LAN port–10G |
| 7 | LAN3 | RJ-45 LAN port–1G |
| 8 | LAN4 | RJ-45 LAN port-1G IPMI-LAN1 |
| 9 | Redundant Power Module | Two 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 offers four 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 3.0 Slots

Skyway is shipped assembled with eight ConnectX-6 VPI cards that enable the hardware-based forwarding of IP packets from InfiniBand to Ethernet, and vice versa.

2.2.4 LAN Interfaces

Skyway features two 10G Base-T Ethernet LAN 1/2 connections (Intel X557 PHY (2* 10GbE)) and two GbE LAN 3/4 (Intel I210 (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 2500 BMC chip that supports IPMI 2.0 (Intelligent Platform Management Interface 2.0) via LAN1 and LAN4 ports.

The recommended cabling when connecting LAN4 and LAN1 (for 10G LAN) is Cat. 6 (< 55m) or Cat. 6A (<100m) for which both cross and straight cables are supported.

LAN3 (X557-AT2 controller) has an internal signal connected to BMC LAN2.

- 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 is equipped with two 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 1+1 redundant power supply
- 1000 W @ 100 ~ 127 V
- 2000 W @ 200 ~ 240 V
- For the redundant power module LEDs definitions, please refer to <u>Power Module LED</u> specifications

2.2.6 Fans Modules

2.2.6.1 Power Supply Fans

NVIDIA Skyway 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 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.

- 2 x 4-pin 80x38 fan for CPU cooling
- 4 x 4-pin 80x38 fan for expansion card cooling
- · All fans are equipped with tachometer status monitoring

3 Hardware Installation

📜 Unknown macro: 'easy-heading-free'

Installation of the NVIDIA Skyway 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 Installation

| Step | Procedure | Direct Link |
|------|---|---------------------------------------|
| 1 | Follow safety warning procedures. | Refer to <u>Safety Warnings</u> |
| 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 <u>Unpacking the Package</u> |
| 5 | Mount the appliance in a rack enclosure. | Refer to Rack Mounting |
| 6 | Use the supplied cables to connect the system | Refer to <u>Cable Installation</u> |
| 7 | Power on the system. | Refer to Initial Power-On |

The installation procedure of NVIDIA Skyway involves the following steps.

3.1.1 Safety Warnings

Prior to the installation, please review the <u>Safety Warnings</u>. 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 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 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 <u>Package Contents</u>.



panel.

3.1.4 Rack Mounting

The NVIDIA Skyway 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 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.



- A 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 (Tma) 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 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 cards, 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

| Left Side IB Port | ConnectX-6 Card | Right Side EN Port | Left Side EN Port | onnectX-6 Card | Right Side IB Port |
|--|-----------------|-----------------------|----------------------|----------------|--|
| HCA1 And And Market And And And And And And | | | | | HCAS Hrms Hubbes HCAS hrms Hubbes |
| HCA3 pume Russ P P HCA4 pume Russ | | | | | HCM Inne Ander |

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 <u>Network Interface Cards LEDs</u> 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 <u>Network Interface</u> <u>Cards LEDs</u> 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 , 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 <u>Technical Specifications</u> 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 <u>System Monitoring</u>.

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).



- a. Go to "Advanced" tab $\rightarrow\,$ "Serial Port Console Redirection" $\rightarrow\,$ Under "Serial Communication via IPMI COM".
- b. 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.



IPMI LAN 1 IPMI LAN 2

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>.</enter> |
| Step 1: Hostname? [gateway-1] | To accept the default hostname, press <enter>. Otherwise, type a different hostname and press <enter>.</enter></enter> |

| Wizard Session Display (Example) | Comments |
|--|--|
| 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></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.</enter> |
| <pre>Step 8: Confirm admin password? <new_password></new_password></pre> | Confirm the password by re-entering it. Note that password characters are not printed. |
| Step 9: Monitor password (Must be typed)? <new_password></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.</enter> |
| Step 10: Confirm monitor password? <new_password></new_password> | Confirm the password by re-entering it. Note that password characters are not printed. |

| Wizard Session Display (Example) | Comments |
|--|---|
| 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] ></gateway></enter></enter></gateway> | 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> Enter the relevant configuration step number, to edit any of the choices To run the command "configuration jumpstart", Config mode must be used. |

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



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.



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".

```
l-csi-skyway-02 login: admin
Password:
Last login: Sun Mar 7 16:05:03 IST 2021 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

Unable to render include or excerpt-include. Could not retrieve page.

5.1 Configuring High Availability (HA)

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

5.1.1 Before Configuring HA

- For all appliances in the HA cluster, the MLNX-GW version must 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 a 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) # exit
eth_router (config) # interface vlan 999 ip address 192.17.10.3/24 primary
eth_router (config) # interface port-channel 1
eth_router (config) # interface port-channel 1
eth_router (config) # interface port-channel 1 pl 1
eth_router (config) # interface port-channel 1 ip 1
eth_router (config) # interface vlan 999 ip 1 peer-address 192.17.10.2
eth_router (config) # interface mlag-vport-channel 101
eth_router (config) # interface mlag-port-channel 101 mode active
eth_router (config) # interface mlag-port-channel 101 no shutdown
eth_router (config) # vlan 101
eth_router (config) # interface vlan 101 ip address 2.2.2.252/24 primary
eth_router (config) # interface vlan 101 ip address 2.2.2.252/24 primary
eth_router (config) # interface vlan 101 ip address 2.2.2.252/24 primary
eth_router (config) # interface vlan 101 ip address 2.2.2.252/24 primary
eth_router (config) # interface vlan 101 imagn 101
eth_router (config) # interface vlan 101 imagn 101
eth_router (config) # interface vlan 101 imagn 101
eth_router (config) # interface vlan 101 magn 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) # 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 1/1-1/18 channel-group 1 mode active
eth_router (config) # vlan 2
eth_router (config) # vlan 2
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.

5.1.2 Configuring HA on Skyway Appliance

4

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.

 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

5.1.2.1 High Availability LAG/MLAG Setup

5.1.2.2 Skyway Connectivity to the Ethernet Using L2 Ethernet Switches

| | n this above use case, every Skyway-facing port on the side of the 12 Fthernet switches should be con | figured |
|---|--|---------|
| Ľ | is an extern part in a difference of the part of the state of the L2 difference of the L2 and the compared of the L2 and the state of t | 1/) |
| İ | as a router port. In addition, a private network should be established (in the example above, $3.3.0.07$ | 10) |
| | between the router ports mentioned above and the Skyways Ethernet port channel. | |
| Ľ | · · · · · · · · · · · · · · · · · · · | |

5.2 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 ALI | , type | full |

3. Configure PKEY on the InfiniBand host.

```
#### pkey_full = pkey_id (hex) + 8000 (hex)
echo "0x<pkey_full>" > /sys/class/net/ibcinterface_number>/create_child
ifconfig ibcinterface_number>.< pkey_full> up
ifconfig ibcinterface_number>.< pkey_full> <lb 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 ib1.8001 up
ifconfig ib1.8001 11.222.62.11 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



Power Module 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 |
| Green | Green | 1G bps linked |
| Green | Blinking Green | 1G 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 | 1G bps linked |
| Amber | Blinking Green | 1G 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 standy 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. |

| LED Color and State | Description |
|------------------------|---|
| 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 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.



| ltem | Description |
|------|-------------------|
| 1 | Power Socket |
| 2 | Extraction handle |
| 3 | PSU Status LEDs |

7.1.1 Extracting and Inserting the Power Supply Unit

A 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.

 \blacktriangleright 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.

 \blacktriangleright 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

0

 \blacktriangleright 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 <u>Slide Rail Installation</u> section.

8 Troubleshooting

As soon as the appliance is plugged in, make sure that the green power LEDs on the power supply units are on.

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 | Check that the power cable is plugged into a working outlet. Check that the power cable has a voltage within the range of 100-240 volts AC. Remove and reinstall the power cable. Remove and reinstall the PSU. |
| The Power Button w/Integrated LED for the appliance shuts off | Check that there is adequate ventilation. 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). 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 appliance has been started. |
| The appliance is off | Press the Power Button w/Integrated LED If that does not work, do the following: Unplug the appliance. Wait 5 minutes. Plug in the appliance, and press the Power Button w/ Integrated LED. If the appliance does not come on, check the power supplies. If the appliance comes on, use the NVIDIA Skyway management software to determine the cause of the shutdown. Check the temperature. Check the fan status. |

9 Technical Specifications

9.1 MGA100-HS2 Specifications

| Physical | Dimensions (W x H x D): 438 x 88 x 760 (17.24" x 3.46" x 29.92") | | | |
|---------------------|--|------------------------------------|--|--|
| | Weight: NVIDIA Skyway gateway: 21kg NVIDIA Skyway gateway with ACC and package: 32kg | | | |
| | Mounting: 19" rack mount | | | |
| 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 | |
| | | Ethernet | 25/50/100/200 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 1+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 | 2x70W | 140W | |
| | 2x PCH | 2x3W | 6W | |
| | 4x DDR4 16G | 4x2W | 8W | |
| | 1x SSD SATA 2.5 64G | 1x2.5W | 2.5W | |
| | 6x fans | 6x52.8W | 316.8W | |
| | 8x ConnectX-6 Network Cards | 8x35W | 280W | |
| | 1x BMC | 1x2W | 2W | |
| | 6x USB 4x RJ45 | 15W 2W | 17W | |
| | Other server components | | 3W | |
| | Maximum Power Consumption of the entire server | | ~774W | |
| Environmental | Temperature | Operating | 0 to 35° C | |
| | | Non-operating | -20 to 60° C | |
| | Humidity | Operating relative humidity | 10~85% @ 40° C (non-condensing) | |

| | | Non-operating relative humidity | 10-95% @ 40 $^{\circ}$ C (non-condensing) |
|-----------------------------|---------------------------|--|---|
| | Vibration (5 ~ 500 Hz) | 0.25 GRMs | |
| | Shock | 10G (with 11ms duration, half | f sine wave) |
| Airflow/Heat Dissipation | Airflow | When all 6 fans are working at full speed: 1045 LFM When only 4 fans are working at full speed: 688 LFM | |

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 which impact the overall system operation state:

 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.

| M Me | llanox | | |
|---|---|---|---|
| <u>近日322(1)</u> FECHNOLOGIES 型号/型號 Model: MGA100 产品名称: 交換机 / 產品名稱: 交換器 Electrical Rating: 1x/2x 100-127Vac, 47/63Hz, 12A 会 複複人 会 入程な、200-240Vac, 47/63Hz, 10A | | | |
| | | | |
| Para Argentina AC 200-246V, 47/63Hz | | | |
| 警告使用者,這是甲類的資源透晶。 此为4级产品。在生活环境中, 在居住的调味中使用時。间密增成。成一局却能会造成无线电干扰。 射頻干擾,在這種情况下,使用者會 在这种构成下,可能需要用户对 被要求某限本是或達的對例。 其干扰某取切实可仿结指施。 | | | |
| European Regulatory Address: Mellanox Technologies Denmar Ledreborg Allé 130B, 4000 Roskilde, Denmark | k ApS. Patented: www.mellanox.com/page/pat | ents | |
| This during a singles with part 15 of the FCC pices Operation is subject to be Labo on litetitist subjectstellmills following two conditions: (11th directions part of cases harmoni interference, received, including and () this device must accept any interference received, including interference that may cause underlead operation. Apparate in all angular states locat states of the states of the states of the states of t | | | |
| CAUTION Risk of electric shock and energy hazard. Disconnecting one power suppl 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 seué alimentation stablisée ne débranche uniquement qu'un module 'Alimentation Stablisée', Pour isoler complétement le module on cause, il faut débrancher toutes les alimentations stablisées. | 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 Gafahr des elektrischen Schocks. Entfermen des Netzsteckers eines Netz -toils spannungsfrei. Um alle Einhieten spannungs-frei zu machen sind die Netzstecker aller Netzteile zu entfernen. |
| DO NOT REMOVE THE COVER www.mellanox.com | | | |



11 Field Replaceable Units

| Ordering Number | Part Description |
|-----------------|---|
| MGA100-PS | Power supply for NVIDIA Skyway Infiniband-to-Ethernet appliance |
| MGA100-RKIT | Rail kit for NVIDIA Skyway Infiniband-to-Ethernet appliance |

12 Revision History

| Date | Description of Changes |
|-----------|---|
| Jun. 2023 | Added the following sections: • <u>Configuring High Availability (HA)</u> • <u>Configuring Partition Keys (PKEYs)</u> |
| Nov. 2022 | Updated Hardware Installation |
| Jun. 2022 | Added partition PKeys to the <u>Networkwide Deployment Guidelines</u> |
| Nov. 2021 | Updated supported protocols across the document. |
| Jun. 2021 | Updated LED tables |
| Feb. 2021 | Updated Technical Specifications. |
| Nov. 2020 | Updated Technical Specifications. |
| Oct. 2020 | Updated package contents. |
| Oct. 2020 | Updated package contents. |
| Aug. 2020 | 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.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, and Mellanox are trademarks and/or registered trademarks of NVIDIA Corporation and/



or Mellanox Technologies Ltd. in the U.S. and in other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2024 NVIDIA Corporation & affiliates. All Rights Reserved.

