



NVIDIA DGX B200 User Guide

NVIDIA Corporation

Dec 09, 2025

Contents

1	Introduction to NVIDIA DGX B200 Systems	3
1.1	Hardware Overview	3
1.1.1	DGX B200 Component Descriptions	3
1.1.2	Mechanical Specifications	4
1.1.3	Power Specifications	5
1.1.3.1	Support for PSU Redundancy and Continuous Operation	5
1.1.4	DGX B200 Locking Power Cord Specification	5
1.1.5	Using the Locking Power Cords	6
1.1.6	Environmental Specifications	7
1.1.7	Front Panel Connections and Controls	7
1.1.7.1	With a Bezel	7
1.1.7.2	With the Bezel Removed	8
1.1.8	Rear Panel Modules	9
1.1.9	Motherboard Connections and Controls	10
1.1.10	Motherboard Tray Components	11
1.1.11	GPU Tray Components	11
1.2	Network Connections, Cables, and Adaptors	12
1.2.1	Network Ports	12
1.2.2	Compute and Storage Networking	13
1.2.3	Network Modules	14
1.2.4	BMC Port LEDs	15
1.2.5	Supported Network Cables and Adaptors	15
1.3	DGX B200 System Topology	16
1.4	DGX OS Software	16
1.5	Customer Support	17
2	Connecting to DGX B200	19
2.1	Connecting to the Console	19
2.1.1	Direct Connection	19
2.1.2	Remote Connection through the BMC	20
2.2	SSH Connection to the OS	22
3	First Boot Setup	23
3.1	System Setup	23
3.2	Post Setup Tasks	25
3.2.1	Obtaining Software Updates	25
3.2.2	Enabling the SRP Daemon	25
4	Quickstart and Basic Operation	27
4.1	Installation and Configuration	27
4.2	Registration	27
4.3	Obtaining an NGC Account	27
4.4	Turning DGX B200 On and Off	28

4.4.1	Startup Considerations	28
4.4.2	Shutdown Considerations	28
4.5	Verifying Functionality - Quick Health Check	28
4.6	Running the Pre-flight Test	29
4.7	Running NGC Containers with GPU Support	29
4.7.1	Using Native GPU Support	30
4.7.2	Using the NVIDIA Container Runtime for Docker	30
4.8	Managing CPU Mitigations	31
4.8.1	Determining the CPU Mitigation State of the DGX System	31
4.8.2	Disabling CPU Mitigations	32
4.8.3	Re-enabling CPU Mitigations	32
5	SBIOS Settings	35
5.1	Accessing the SBIOS Setup	35
5.2	Configuring the Boot Order	36
5.3	Configuring the Local Terminal	37
5.3.1	Linux	38
5.3.2	Windows and MacOS	38
5.4	Power on or Reboot the System	38
6	Using the Baseboard Management Controller (BMC)	39
6.1	Connecting to the BMC	39
6.2	Overview of BMC Controls	40
6.3	Open Ports	43
6.4	Configuring a Static IP Address for the BMC	43
6.4.1	Configuring a BMC Static Address by Using ipmitool	43
6.4.2	Configuring a BMC Static IP Address by Using the System BIOS	44
6.5	Changing the BMC Login Credentials	44
6.5.1	Username and Password Requirements	44
6.5.2	Procedure	45
6.6	Using the Remote Console	45
6.7	Setting Up Active Directory, LDAP, or E-Directory	45
6.8	Configuring Platform Event Filters	46
6.9	Uploading or Generating SSL Certificates	47
6.9.1	Viewing the SSL Certificate	47
6.9.2	Generating the SSL Certificate	47
6.9.3	Uploading the SSL Certificate	48
6.9.4	Updating the SBIOS Certificate	49
7	Managing Power Capping	53
7.1	Querying the Current GPU Power Limit	53
7.2	Managing N+N Configuration (IPMI)	54
7.3	Managing Power Capping Using Redfish API	54
8	Security	55
8.1	User Security Measures	55
8.1.1	Securing the BMC Port	55
8.2	System Security Measures	55
8.2.1	Secure Flash of DGX B200 Firmware	55
8.2.2	Encryption	55
8.2.3	NVIDIA System Manager Security	56
8.3	Secure Data Deletion	56
8.3.1	Prerequisites	56
8.3.2	Procedure	56


9 Redfish APIs Support	59
9.1 Supported Redfish Features	59
9.2 Connectivity Between the Host and BMC	60
9.3 Redfish Examples	60
9.3.1 BMC Manager	60
9.3.2 Firmware Update	61
9.3.3 BIOS Settings	63
9.3.4 Modifying the Boot Order Using Redfish	65
9.3.5 Telemetry	68
9.3.6 Chassis	68
9.3.7 SEL Logs	69
9.3.8 Virtual Image	69
9.3.9 Backing Up and Restoring BMC Configurations	72
9.3.9.1 Backing Up the BMC Configuration	73
9.3.9.2 Restoring the BMC Configuration	73
9.3.10 Collecting BMC Debug Data	73
9.3.11 Clear BIOS and Reset to Factory Defaults	75
9.3.12 Querying GPU Power Limit	75
9.3.13 Power Capping	76
9.3.13.1 Services	76
9.3.13.2 Domains	76
9.3.13.3 Custom Policies	79
9.3.13.4 PSU Policies	82
10 Safety	89
10.1 Safety Information	89
10.2 Safety Warnings and Cautions	89
10.3 Intended Application Uses	91
10.4 Site Selection	91
10.5 Equipment Handling Practices	92
10.6 Electrical Precautions	92
10.6.1 Power and Electrical Warnings	92
10.6.2 Power Cord Warnings	93
10.7 System Access Warnings	93
10.8 Rack Mount Warnings	94
10.9 Electrostatic Discharge	94
10.10 Other Hazards	95
10.10.1 CALIFORNIA DEPARTMENT OF TOXIC SUBSTANCES CONTROL	95
10.10.2 NICKEL	95
10.10.3 Battery Replacement	95
10.10.4 Cooling and Airflow	96
11 Compliance	97
11.1 United States	97
11.2 United States/Canada	97
11.3 Canada	98
11.4 CE	98
11.5 Australia and New Zealand	98
11.6 Brazil	99
11.7 Japan	99
11.8 South Korea	101
11.9 China	101
11.10 Taiwan	102
11.11 Russia/Kazakhstan/Belarus	104

11.12	Israel	104
11.13	India	104
11.14	South Africa	105
11.15	Great Britain (England, Wales, and Scotland)	105
12	Third-Party License Notices	107
12.1	Micron msecli	107
12.2	Mellanox (OFED)	108
13	Notices	109
13.1	Notice	109
13.2	Trademarks	110

The *NVIDIA DGX B200 System User Guide* is also available as a [PDF](#).

Chapter 1. Introduction to NVIDIA DGX B200 Systems

The NVIDIA DGX™ B200 System is the universal system purpose-built for all AI infrastructure and workloads from analytics to training to inference. The system is built on eight NVIDIA B200 Tensor Core GPUs.



1.1. Hardware Overview

1.1.1. DGX B200 Component Descriptions

The NVIDIA DGX B200 (1,440 GB) system includes the following components.

Table 1: Component Description

Component	Description
GPU	8 x NVIDIA B200 GPUs that provide 1,440 GB total GPU memory
CPU	2 x Intel Xeon 8570 PCIe Gen5 CPUs with 56 cores each 2.1/4 GHz (Base/Max boost)
NVSwitch	2 x 5th generation NVLink switches that provide 14.4 TB/s aggregate bandwidth
Storage (OS)	2 x 1.92 TB NVMe M.2 SSD (ea) in RAID 1 array
Storage (Data Cache)	8 x 3.84 TB NVMe U.2 SED (ea) in RAID 0 array
Network (Cluster) card	4 x OSFP ports for 8 x NVIDIA® ConnectX®-7 Single Port Cards Each card provides the following speeds: <ul style="list-style-type: none">▶ InfiniBand (default): Up to 400Gbps▶ Ethernet: 400GbE, 200GbE, 100GbE, 50GbE, 40GbE, 25GbE, and 10GbE
Network (storage and in-band management) card	2 x NVIDIA® BlueField®-3 DPU Dual Port Cards Each card provides the following speeds: <ul style="list-style-type: none">▶ Ethernet (1 port): 400GbE, 200GbE, 100GbE, 50GbE, 40GbE, 25GbE, and 10GbE▶ InfiniBand (1 port): Up to 400Gbps
System memory (DIMM)	2 TB using 32 x DIMMs (upgradable to 4 TB)
BMC (out-of-band system management)	1 GbE RJ45 interface Supports Redfish, IPMI, SNMP, KVM, and Web user interface
System management interfaces (optional)	Dual port 100GbE in slot 3 and 10 GbE RJ45 interface
Power supply	6 x 3.3 kW

1.1.2. Mechanical Specifications

Table 2: Mechanical Specifications

Feature	Description
Form Factor	10U Rackmount
Height	17.5" (444 mm)
Width	19" (482.3 mm) max
Depth	35.3" (897.1 mm) max
System Weight	313.9 lbs (142.4 kg) max

1.1.3. Power Specifications

The DGX B200 system contains six power supplies with a balanced distribution of the power load.

Table 3: Power Specifications

Input	Specification for Each Power Supply	
200-240 volts AC	14.3 kW max.	3,300 W @ 200-240 V, 16 A, 50-60 Hz

1.1.3.1 Support for PSU Redundancy and Continuous Operation

The system includes six power supply units (PSU) configured for 5+1 redundancy.

Refer to the following additional considerations:

- ▶ If a PSU fails, troubleshoot the cause and replace the failed PSU immediately.
- ▶ If two PSUs lose power due to a data center issue or power distribution unit failure, the system continues to function but at a reduced performance level.
- ▶ If only two PSUs have power, shut down the system before replacing an operational PSU.
- ▶ The system only boots if at least three PSUs are operational. If fewer than three are operational, only the BMC is available.
- ▶ Do not operate the system with PSUs depopulated.

This table shows the effects of disconnecting AC power cables from the DGX B200 system.

Table 4: System Effects of Disconnecting AC Power Cables

AC Power Cable Disconnection	Operating PSUs	Effect	Maximum Power per GPU ¹
Removing 1 cable	5	No effect on the DGX system.	1,000 W
Removing 2 cables	4	The BMC alerts on AC power loss.	800 W
Removing 3 cables	3	The BMC alerts on AC power loss.	800 W
Removing 4 cables	2	The DGX system shuts down.	Not applicable

1.1.4. DGX B200 Locking Power Cord Specification

The DGX B200 system is shipped with a set of six (6) locking power cords that have been qualified for use with the DGX B200 system to ensure regulatory compliance.

¹ The maximum power per GPU is reported by the nvidia-smi tool.

⚠ Warning

To avoid electric shock or fire, only use the NVIDIA-provided power cords to connect power to the DGX B200. For more information, refer to [Electrical Precautions](#).

⚠ Important

Do not use the provided cables with any other product or for any other purpose.

Power Cord Specification

Power Cord Feature	Specification
Electrical	250VAC, 20A
Plug Standard	C19/C20
Dimension	1200mm length
Compliance	Cord: UL62, IEC60227 Connector/Plug: IEC60320-1


1.1.5. Using the Locking Power Cords

This section provides information about how to use the locking power cords.

Locking and Unlocking the PDU Side

Power Distribution Unit side

- ▶ To INSERT, push the cable into the PDU socket.
- ▶ To REMOVE, press the clips together and pull the cord out of the socket.



Locking/Unlocking the PSU Side (Cords with Twist-Lock Mechanism)

Power Supply (System) side - Twist locking

- ▶ To INSERT or REMOVE, ensure the cable is UNLOCKED and push/ pull into/out of the socket.



To UNLOCK the power cord, twist the gray locking ring to the unlocked (indicator will show an unlocked padlock)



To LOCK the power cord, twist the gray locking ring to the locked position (indicator should show a locked padlock)

1.1.6. Environmental Specifications

Here are the environmental specifications for your DGX B200 system.


Feature	Specification
Operating Temperature	10°C to 35°C (50°F to 95°F)
Relative Humidity	20% to 80% non-condensing
Airflow	1,550 CFM
Heat Output	48,794 BTU/hr

1.1.7. Front Panel Connections and Controls

This section provides information about the front panel, connections, and controls of the DGX B200 system.

1.1.7.1 With a Bezel


Here is an image of the DGX B200 system with a bezel.



Control	Description
Power Button	<p>Press to turn the DGX B200 system on or off.</p> <ul style="list-style-type: none"> ▶ Green flashing (1 Hz): Standby (BMC booted) ▶ Green flashing (4 Hz): POST in progress ▶ Green solid On: Power On
ID Button	<p>Press to have the blue LED turn On or blink (configurable through the BMC) as an identifier during servicing.</p> <p>It also causes an LED on the back of the unit to flash as an identifier during servicing.</p>
Fault LED	Amber On: System or component faulted

1.1.7.2 With the Bezel Removed

Here is an image of the DGX B200 system without a bezel.




Important

Refer to the section [First Boot Setup](#) for instructions on how to properly turn the system on or off.

1.1.8. Rear Panel Modules

Here is an image that shows the actual panel modules on DGX B200.



1.1.9. Motherboard Connections and Controls

The following image shows the motherboard connections and controls in a DGX B200 system.





Table 5: Motherboard Controls

Control	Description	
Power Button	Press to turn the system On or Off.	
ID LED Button	It blinks when the ID button is pressed from the front of the unit to help identify the unit that needs servicing.	
BMC Reset Button	Press to manually reset the BMC.	

See [Network Connections, Cables, and Adapters](#) for details on the network connections.


1.1.10. Motherboard Tray Components

The following image shows the motherboard tray components in the DGX B200 system.



1.1.11. GPU Tray Components

Here is an image of the GPU tray components in the DGX B200 system.



1.2. Network Connections, Cables, and Adaptors

This section provides information about network connections, cables, and adaptors.

1.2.1. Network Ports

Here is an image that shows the network ports on a DGX B200 system.





Table 6: Network Port Mapping

Port Designation				
Port	PCI Bus	Default	Optional	RDMA
OSFP1P1	dc:00.0	ibp220s0	enp220s0np0	mlx5_15
OSFP1P2	9a:00.0	ibp154s0	enp154s0np0	mlx5_10
OSFP2P1	ce:00.0	ibp206s0	enp206s0np0	mlx5_14
OSFP2P2	c0:00.0	ibp192s0	enp192s0np0	mlx5_13
OSFP3P1	4f:00.0	ibp79s0	enp79s0np0	mlx5_8
OSFP3P2	40:00.0	ibp64s0	enp64s0np0	mlx5_7
OSFP4P1	5e:00.0	ibp94s0	enp94s0np0	mlx5_9
OSFP4P2	18:00.0	ibp24s0	enp24s0np0	mlx5_4
Slot1 P1	aa:00.0	ibp170s0f0	enp170s0f0np0	mlx5_11
Slot1 P2	aa:00.1	enp170s0f1np1	ibp170s0f1	mlx5_12
Slot2 P1	29:00.0	ibp41s0f0	enp41s0f0np0	mlx5_5
Slot2 P2	29:00.1	enp41s0f1np1	ibp41s0f1	mlx5_6
Slot3 P1	82:00.0	ens6f0	N/A	irdma0
Slot3 P2	82:00.1	ens6f1	N/A	irdma1
On-board	0b:00.0	eno3	N/A	

1.2.2. Compute and Storage Networking




1.2.3. Network Modules

- ▶ New form factor for aggregate PCIe network devices
- ▶ Consolidates four ConnectX-7 networking cards into a single device

The DGX B200 system has eight ConnectX-7 network cards on two network module trays. Internal DensiLink cables connect the dual-port OSFP interface to the individual ConnectX-7 network card.

Table 7: Network Modules

Port	ConnectX Device	Network Module/CPU	GPU	Default	RDMA
OSFP1P1	CX0	1	7	ibp220s0	mlx5_15
OSFP1P2	CX1	1	4	ibp154s0	mlx5_10
OSFP2P1	CX2	1	6	ibp206s0	mlx5_14
OSFP2P2	CX3	1	5	ibp192s0	mlx5_13
OSFP3P1	CX2	0	2	ibp79s0	mlx5_8
OSFP3P2	CX3	0	1	ibp64s0	mlx5_7
OSFP4P1	CX0	0	3	ibp94s0	mlx5_9
OSFP4P2	CX1	0	0	ibp24s0	mlx5_4



1.2.4. BMC Port LEDs

The BMC RJ-45 port has two LEDs.

The LED on the left indicates the speed. Solid green indicates the speed is 100M. Solid amber indicates the speed is 1G.

The LED on the right is green and flashes to indicate activity.

1.2.5. Supported Network Cables and Adaptors

The DGX B200 system is not shipped with network cables or adaptors. You will need to purchase supported cables or adaptors for your network.

The ConnectX-7 firmware determines which cables and adaptors are supported. For a list of cables and adaptors compatible with the NVIDIA ConnectX cards installed in the DGX B200 system,


1. Visit the [NVIDIA Adapter Firmware Release page](#).
2. Click the ConnectX model and select the corresponding firmware included in the DGX B200 system.

3. From the left **Topics** pane, select the Validated and Supported Cables and Switches topic.

To configure the BlueField-3 DPU in NIC mode, follow the instructions in [NIC Mode for BlueField-3](#).

1.3. DGX B200 System Topology

The following figure shows the DGX B200 system topology.



1.4. DGX OS Software

The DGX B200 system comes pre-installed with a DGX software stack incorporating the following components:

- ▶ An Ubuntu server distribution using the optimized Linux kernel with supporting packages
- ▶ The following system management and monitoring software:
 - ▶ NVIDIA System Management (NVSM)

Provides active health monitoring and system alerts for NVIDIA DGX nodes in a data center. It also provides simple commands for checking the health of the DGX B200 system from the command line.
 - ▶ Data Center GPU Management (DCGM)

This software enables node-wide administration of GPUs and can be used for cluster and data-center level management.
 - ▶ DGX B200 system support packages
 - ▶ The NVIDIA GPU driver, including NVIDIA CUDA

- ▶ Docker Engine
- ▶ NVIDIA Container Toolkit
- ▶ NVIDIA Networking OpenFabrics Enterprise Distribution for Linux (DOCA-OFED)
- ▶ NVIDIA Networking Software Tools (MST)
- ▶ cachefilesd (daemon for managing cache data storage)

1.5. Customer Support

Contact NVIDIA Enterprise Support for assistance in reporting, troubleshooting, or diagnosing problems with your DGX B200 system. You can also contact NVIDIA Enterprise Support for help in moving the DGX B200 system.

- ▶ For contracted Enterprise Support questions, you can send an email to enterprisesupport@nvidia.com.
- ▶ For more information on obtaining support, go to [NVIDIA Enterprise Support](#).

Our support team can help collect appropriate information about your issue and involve internal resources as needed.

Chapter 2. Connecting to DGX B200

2.1. Connecting to the Console

Connect to the DGX B200 console using either a direct connection or a remote connection through the BMC.

Important

Connect directly to the DGX B200 console if the NVIDIA DGX™ B200 system is connected to a 172.17.xx.xx subnet.

DGX OS Server software installs Docker Engine, which uses the 172.17.xx.xx subnet by default for Docker containers. If the DGX B200 system is on the same subnet, you cannot establish a network connection to the DGX B200 system.

For instructions on changing the default Docker network settings, refer to [Configuring Docker IP Addresses](#) in the *NVIDIA DGX OS 7 User Guide*.


2.1.1. Direct Connection

At the front or the back of the system, you can connect a display to the VGA connector and a keyboard to any of the USB ports.

The system provides video to one of the two VGA ports at a time. Simultaneous video output is not supported. If you connect both VGA ports, the VGA port on the rear has precedence.

Note

The display resolution must be 1440x900 or lower.



2.1.2. Remote Connection through the BMC

Here is some information about remotely connecting to DGX B200 through the BMC.

NVIDIA recommends that customers follow best security practices for BMC management (IPMI port). These include, but are not limited to, such measures as:

- ▶ Restricting the DGX B200 IPMI port to an isolated, dedicated management network.
- ▶ Using a separate, firewalled subnet.
- ▶ Configuring a separate VLAN for BMC traffic if a dedicated network is unavailable.

This method requires that you have the BMC login credentials. These credentials depend on the following conditions:

Before the First Boot Setup

 **Caution**

Perform the first boot setup to change the default credentials before you connect the BMC to an unsecured network.

- ▶ The default credentials are:

- ▶ Username: admin
- ▶ Password: admin

 **Caution**

When you create a BMC admin user, we strongly recommend changing the default password for this user - DO NOT use the default password.

After the First Boot Setup

During the first boot procedure, you were prompted to configure an administrator username and password and a password for the BMC. The BMC username is the same as the administrator username:


- ▶ Username: <administrator-username>
- ▶ Password: <bmc-password>

1. Ensure you have connected the BMC port on the DGX B200 system to your LAN.
2. Open a browser within your LAN and go to <https://<bmc-ip-address>/>
Ensure popups are allowed for the BMC address.
3. Log in.
4. From the navigation menu, click **Remote Control**.

The **Remote Control** page lets you open a virtual Keyboard/Video/Mouse (KVM) on the DGX B200 system just like you are using a physical monitor and keyboard connected to the front of the system.

5. Click **Launch KVM**.

The DGX B200 console appears in your browser.



2.2. SSH Connection to the OS

After configuring the system, you can establish an SSH connection to the DGX B200 OS through the network port. Refer to [Network Ports](#) to identify the port to use.

Chapter 3. First Boot Setup

This section provides information about the setup process after you first boot the NVIDIA DGX™ B200 system.


While NVIDIA partner network personnel or NVIDIA field service engineers will install the DGX B200 system at the site and perform the first boot setup, the first boot setup instructions are provided here for reference and to support any reimaging of the server.

3.1. System Setup


These instructions describe the setup process that occurs the first time the DGX B200 system is powered on after delivery or after the server is re-imaged.

Be prepared to accept all End User License Agreements (EULAs) and to set up your username and password. To preview the EULA, visit the [Enterprise Software Agreements](#) page and click [NVIDIA Software License Agreement](#) under **NVIDIA AI, DGX and NVIDIA AI Enterprise**.

1. Connect to the DGX B200 console as explained in [Connecting to the Console](#).
2. Power on the DGX B200 system in one of the following ways:
 - ▶ Using the physical power button.



► Using the Remote BMC



3. Refer to [First Boot Process for DGX Servers](#) in the *NVIDIA DGX OS 7 User Guide* for information about the following topics:
 - ▶ Optionally encrypt the root file system.
 - ▶ Use the first boot wizard to set the language, locale, country, and so on.
 - ▶ Create an administrative user account for the system, BMC, and Grub boot loader.

- ▶ Configure the primary network interface.

3.2. Post Setup Tasks

This section explains recommended tasks to perform after the initial system first-boot setup.

Note

RAID 1 rebuild can temporarily affect system performance.

When the system is booted after restoring the image and running the first-boot setup, software RAID begins rebuilding the RAID 1 array, which creates a mirror of (or resynchronizes) the drive containing the software. System performance can be affected during the RAID 1 rebuild process. The process can take an hour to complete.

During this time, running the `nvsm show health` command reports a warning that the RAID volume is re-syncing.

You can monitor the status of the RAID 1 rebuild process by running the `sudo nvsm show volumes` command and then view the output under `/systems/localhost/storage/volumes/md0/rebuild`.

3.2.1. Obtaining Software Updates

To ensure that you are running the latest version of DGX OS, you might need to update the software.

Updating the software ensures that your DGX B200 system contains important updates, including security updates. The Ubuntu Security Notice site, <https://usn.ubuntu.com/>, lists known common vulnerabilities and exposures (CVEs), including those that can be resolved by updating the DGX OS software.

Refer to [Upgrading the OS](#) in the *NVIDIA DGX OS 7 User Guide* for information about updating the operating system.

3.2.2. Enabling the SRP Daemon

The NVIDIA networking drivers provide the SRP daemon software. The daemon is disabled by default. Enabling the daemon is required if you want to use RDMA over InfiniBand. You can enable the daemon by running the following commands:

```
sudo systemctl enable srp_daemon.service
sudo systemctl enable srptools.service
```

Chapter 4. Quickstart and Basic Operation

This topic provides basic requirements and instructions for using the NVIDIA DGX™ B200 system, including performing a preliminary health check and preparing for running containers. For additional product documentation, refer to the [NVIDIA DGX Platform](#) page.

4.1. Installation and Configuration

Before you install DGX B200, ensure you have given all relevant site information to your Installation Partner.

 **Important**

Your DGX B200 system must be installed by NVIDIA partner network personnel or NVIDIA field service engineers. If not performed accordingly, your hardware warranty will be voided.

4.2. Registration

To obtain support for your DGX B200 system, follow the registration instructions in the Entitlement Certification email that was sent with the purchase.

Registration allows you to access the NVIDIA Enterprise Support Portal, obtain technical support, get software updates, and set up an NGC for DGX systems account. If you did not receive the information, open a case with the NVIDIA Enterprise Support Team at [Enterprise Support Services](#).

Refer to [Customer Support](#) for contact information.

4.3. Obtaining an NGC Account

NVIDIA NGC provides access to GPU-optimized software for deep learning, machine learning, and high-performance computing (HPC). An NGC account grants you access to these tools and allows you to set up a private registry to manage your customized software.

If you are the organization administrator for your DGX system purchase, work with NVIDIA Enterprise Support to set up an NGC enterprise account. Refer to the [NGC Private Registry User Guide](#) for more information about getting an NGC enterprise account.

4.4. Turning DGX B200 On and Off

DGX B200 is a complex system that integrates a large number of cutting-edge components with specific startup and shutdown sequences. Observe the following startup and shutdown instructions.

4.4.1. Startup Considerations

To keep your DGX B200 running smoothly, allow up to a minute of idle time after reaching the login prompt. This ensures that all components can complete their initialization.

4.4.2. Shutdown Considerations

When shutting down DGX B200, always initiate the shutdown from the operating system, momentary press of the power button, or by using Graceful Shutdown from the BMC, and wait until the system enters a powered-off state before performing any maintenance.

⚠ Warning

Risk of Danger - Removing power cables or using Power Distribution Units (PDUs) to shut off the system while the Operating System is running may cause damage to sensitive components in the DGX B200 server.

4.5. Verifying Functionality - Quick Health Check

NVIDIA provides customers a diagnostics and management tool called NVIDIA System Management, or NVSM. The `nvsm` command can be used to determine the system's health, identify component issues and alerts, or run a stress test to ensure all components are in working order while under load. The use of Docker is key to getting the most performance out of the system since NVIDIA has optimized containers for all the major frameworks and workloads used on DGX systems.

The following instructions show how to perform a health check on the DGX B200 system and verify the Docker and NVIDIA driver installation.

1. Establish an SSH connection to the DGX B200 system.
2. Run a basic system check.

```
sudo nvsm show health
```

3. Verify that the output summary shows that all checks are Healthy and that the overall system status is Healthy.
4. Verify that Docker is installed by viewing the installed Docker version.

```
sudo docker --version
```

On success, the command returns the version as `Docker version xx.yy.zz`, where the actual version may differ depending on the specific release of the DGX OS Server software.

5. Verify connection to the NVIDIA repository and that the NVIDIA Driver is installed.

```
sudo docker run --gpus all --rm nvcr.io/nvidia/cuda:12.1.1-base-ubuntu22.  
    ↪04 nvidia-smi
```

The preceding command pulls the nvidia/cuda container image layer by layer, then runs the nvidia-smi command.

When complete, the output shows the NVIDIA Driver version and a description of each installed GPU.

For more information, refer to *Containers For Deep Learning Frameworks User Guide*.

4.6. Running the Pre-flight Test

Instructions for running the DGX stress test.

NVIDIA recommends running the pre-flight stress test before putting a system into a production environment or after servicing. You can specify running the test on the GPUs, CPU, memory, and storage, and also specify the duration of the tests.

To run the tests, use NVSM.

Syntax

```
sudo nvsm stress-test [--usage] [--force] [--no-prompt] [<test>...] [DURATION]
```

For help on running the test, issue the following.

```
sudo nvsm stress-test --usage
```

Recommended Command

The following command runs the test on all supported components (GPU, CPU, memory, and storage), and takes approximately 20 minutes.

```
sudo nvsm stress-test --force
```

4.7. Running NGC Containers with GPU Support

To obtain the best performance when running NGC containers on the DGX B200 system, the following methods of providing GPU support for Docker containers are available:

- ▶ Native GPU support (included in Docker 20.10.18 and later)

The method implemented in your system depends on the DGX OS version installed.

DGX OS Releases	Method Included
7.0	<ul style="list-style-type: none"> ▶ Native GPU support ▶ NVIDIA Container Runtime for Docker (deprecated - availability to be removed in a future DGX OS release)

Each method is invoked by using specific Docker commands, described as follows.

4.7.1. Using Native GPU Support

Use `docker run --gpus` to run GPU-enabled containers.

- ▶ Example using all GPUs

```
sudo docker run --gpus all ...
```

- ▶ Example using two GPUs

```
sudo docker run --gpus 2 ...
```

- ▶ Examples using specific GPUs

```
sudo docker run --gpus '{"device=1,2"' ...  
sudo docker run --gpus '{"device=UUID-ABCDEF,1"' ...
```

4.7.2. Using the NVIDIA Container Runtime for Docker

If you need to use `nvidia-docker2`, install it using `sudo apt install nvidia-docker2`, then run:

```
sudo systemctl restart docker
```

The DGX OS also includes the NVIDIA Container Runtime for Docker (`nvidia-docker2`), which lets you run GPU-accelerated containers in one of the following ways:

- ▶ Use `docker run` and specify `runtime=nvidia`.

```
docker run --runtime=nvidia ...
```

- ▶ Use `nvidia-docker run`.

```
nvidia-docker run ...
```

The `nvidia-docker2` package provides backward compatibility with the previous `nvidia-docker` package, so you can run GPU-accelerated containers using this command and the new runtime will be used.

- ▶ Use `docker run` with `nvidia` as the default runtime.

You can set `nvidia` as the default runtime, for example, by adding the following line to the `/etc/docker/daemon.json` configuration file as the first entry.

```
"default-runtime": "nvidia",
```

Here is an example of how the added line appears in the JSON file. Do not remove any pre-existing content when making this change.

```
{
  "default-runtime": "nvidia",
  "runtimes": {
    "nvidia": {
      "path": "/usr/bin/nvidia-container-runtime",
      "args": []
    }
  }
}
```

You can then use docker run to run GPU-accelerated containers.

```
docker run ...
```

⚠️ Caution

If you build Docker images while nvidia is set as the default runtime, ensure the build scripts executed by the Dockerfile specify the GPU architectures that the container will need. Failure to do so might result in the container being optimized only for the GPU architecture on which it was built. Instructions for specifying the GPU architecture depend on the application and are beyond the scope of this document. Consult the specific application build process.

For more information, refer to the [NVIDIA DGX OS 7 User Guide](#).

4.8. Managing CPU Mitigations

DGX OS Server includes security updates to mitigate CPU speculative side-channel vulnerabilities. These mitigations can decrease the performance of deep learning and machine learning workloads.

If your installation of DGX systems incorporates other measures to mitigate these vulnerabilities, such as measures at the cluster level, you can disable the CPU mitigations for individual DGX nodes and thereby increase performance. This capability is available starting with DGX OS Server release 4.4.0.

4.8.1. Determining the CPU Mitigation State of the DGX System

If you do not know whether CPU mitigations are enabled or disabled, issue the following.

```
cat /sys/devices/system/cpu/vulnerabilities/*
```

- ▶ CPU mitigations are enabled if the output consists of multiple lines prefixed with **Mitigation:**.

Example

```

KVM: Mitigation: Split huge pages
Mitigation: PTE Inversion; VMX: conditional cache flushes, SMT vulnerable
Mitigation: Clear CPU buffers; SMT vulnerable
Mitigation: PTI
Mitigation: Speculative Store Bypass disabled via prctl and seccomp
Mitigation: usercopy/swapgs barriers and __user pointer sanitization
Mitigation: Full generic retpoline, IBPB: conditional, IBRS_FW, STIBP:
  ↪conditional, RSB filling
Mitigation: Clear CPU buffers; SMT vulnerable

```

- CPU mitigations are disabled if the output consists of multiple lines prefixed with Vulnerable.

Example

```

KVM: Vulnerable
Mitigation: PTE Inversion; VMX: vulnerable
Vulnerable; SMT vulnerable
Vulnerable
Vulnerable
Vulnerable: __user pointer sanitization and usercopy barriers only; no swapgs
  ↪barriers
Vulnerable, IBPB: disabled, STIBP: disabled
Vulnerable

```

4.8.2. Disabling CPU Mitigations

⚠ Caution

Performing the following instructions will disable the CPU mitigations provided by the DGX OS Server software.

1. Install the nv-mitigations-off package.

```
sudo apt install nv-mitigations-off -y
```

2. Reboot the system.

3. Verify CPU mitigations are disabled.

```
cat /sys/devices/system/cpu/vulnerabilities/*
```

The output should include several Vulnerable lines. For an example output, refer to [Determining the CPU Mitigation State of the DGX System](#).

4.8.3. Re-enabling CPU Mitigations

1. Remove the nv-mitigations-off package.

```
sudo apt purge nv-mitigations-off
```

2. Reboot the system.

3. Verify CPU mitigations are enabled.

```
cat /sys/devices/system/cpu/vulnerabilities/*
```

The output should include several **Mitigations** lines. See [Determining the CPU Mitigation State of the DGX System](#) for example output.

Chapter 5. SBIOS Settings

The NVIDIA DGX™ B200 system contains a system BIOS with optimized settings for the DGX system. There might be situations where the settings need to be changed, such as changes in the boot order, changes to enable PXE booting, or changes in the BMC network settings.

Instructions for these use cases are provided in this section.

 **Important**

Do not change settings in the SBIOS other than those described in this or other DGX B200 user documents. Contact NVIDIA Enterprise Services **before** making other changes.

5.1. Accessing the SBIOS Setup

Here is information about how you can access the SBIOS setup.

1. Access the DGX B200 console from a locally connected keyboard and mouse or through the BMC remote console.
2. Reboot the DGX B200 system.
3. When presented with the SBIOS version screen, press the Del or F2 key to enter the BIOS Setup Utility.




Here are some occasions where it might be necessary to reconfigure settings in the SBIOS:

- ▶ Configuring a BMC static IP address using the system BIOS
- ▶ Enabling the TPM and preventing the BIOS from sending Block SID requests
For more information, refer to [Enabling the TPM and Preventing the BIOS from Sending Block SID Requests](#).
- ▶ Clearing the TPM

5.2. Configuring the Boot Order


The following instructions describe how to set the boot order at boot time. You can also set the boot order from the SBIOS setup > Boot screen.

1. Access the DGX B200 console from a locally connected keyboard and mouse or through the BMC remote console.
2. Reboot the DGX B200 system.
3. Press the F11 key on the NVIDIA splash screen.



4. Select the boot device.

The following figure shows the virtual media selected.



5.3. Configuring the Local Terminal

There are two ways to access the BIOS setup screen:

- ▶ A direct-attached keyboard and monitor
- ▶ Serial-over-LAN (SOL) using SSH or IPMItool

To use the SOL connection, you might need to configure your terminal application.

5.3.1. Linux

1. Set the locale and language for your terminal:

```
sudo localectl set-locale LANG=en_US.UTF-8
```

2. Set the locale for the current session:

```
export LANG=en.UTF-8
```

3. Type `xterm` to launch the terminal with the set locale.

5.3.2. Windows and MacOS

- ▶ Configure your terminal application for `en_US.UTF-8` support.

5.4. Power on or Reboot the System

1. Reboot the system using one of the following methods:

- ▶ Connect to the BMC web interface and click **power on/reboot**.
- ▶ From an operating system command line, run `sudo reboot`.

2. Connect to the DGX B200 SOL console:

- ▶ Using SSH

1. Create a new user using the BMC Web UI, for example, `userA`.
2. Enable the **solssh** service on the **Services** page of the Web UI.
3. Restart the SSH service and log in as the new user created in step A.

For example,

```
ssh userA@10.33.128.52
```

- ▶ Using IPMItool

```
ipmitool -I lanplus -H <ip-address> -U admin -P dgxluna.admin sol
  ↵activate
```

3. Press the **Del** or **F2** key when the system is booting.

The system confirms your choice and shows the BIOS configuration screen.

Chapter 6. Using the Baseboard Management Controller (BMC)

The NVIDIA DGX™ B200 system contains a baseboard management controller (BMC) that monitors and controls various hardware devices on the system, including system sensors and other parameters.

6.1. Connecting to the BMC

Here are the steps to connect to the BMC on a DGX B200 system.

Before you begin, ensure that you connect the BMC network interface controller port on the DGX system to your LAN.


1. Open a browser within your LAN and enter the IP address of the BMC in the location.

The BMC is supported on the following browsers:

- ▶ Internet Explorer 11 and later
- ▶ Firefox 29.0 (64-bit) and later
- ▶ Google Chrome 7.0.3396.87 (64-bit) and later

2. Log in.

The BMC dashboard opens.



6.2. Overview of BMC Controls

The left-side navigation menu bar on the BMC main page contains the primary controls.

NVIDIA DGX™ B200

UpTime : 0 d 1 hr

IP :

MAC : 5C:FF:35:FC:03:E9

Chassis PN : 675-24387-0061-000

Chassis SN : 1663724000066

● Host Online

● Chassis Identify LED

Quick Links.. ▾

- Dashboard
- Sensor
- System Inventory
- FRU Information
- GPU Information
- Firmware Inventory
- Logs & Reports >
- Settings
- Remote Control
- Image Redirection
- Host System Diagnostics
- Power Control
- Chassis ID LED Control
- Maintenance
- Sign out

Table 1: BMC Main Controls

Control	Description
Quick Links	Provide quick access to several tasks.
Dashboard	Display the overall information about the status of the device.
Sensor	Provide status and readings for system sensors, such as SSD, PSUs, voltages, CPU temperatures, DIMM temperatures, and fan speeds.
System Inventory	Display inventory information of system modules.
FRU Information	Provide system, processor, memory controller, baseboard, power, thermal, PCIE device, PCIE function, and storage.
GPU Information	Provide basic information on all the GPUs in the systems, including GUID, VBIOS version, InfoROM version, and number of retired pages for each GPU.
Logs and Reports	View, and if applicable, download and erase, the IPMI event log, and System, Audit, Video, and POST Code logs.
Settings	Configure the following settings: Captured BSOD, External User Services, KVM Mouse Settings, Log Settings, Media Redirection Settings, Network Settings, PAM Order Settings, Platform Event Filter, Services, SMTP Settings, SSL Settings, System Firewall, User Management, and Video Recording
Remote Control	Open the KVM Launch page to access the DGX B200 console remotely.
Power Control	Perform the following power actions: Power On, Power Off, Power Cycle, Hard Reset, and ACP/Shutdown
Chassis ID LED Control	Virtual LED is a button to toggle the UID LED on/off: <ul style="list-style-type: none"> ▶ Off ▶ Solid on ▶ Blinking on (select from 5 to 255-second blink interval). This is activated by the Chassis Identify LED option above the Quick Links dropdown.
Maintenance	Perform the following maintenance tasks: Backup Configuration, Firmware Image Location, Firmware Update, Preserve Configuration, Restore Configuration, Restore Factory Defaults, and System Administrator
Sign out	Sign out of the BMC web UI.

6.3. Open Ports

Ensure that the ports listed in the following table are open and available on your firewall to the DGX B200 system.

Table 2: Open Ports

Port	Protocol	Function
443	HTTPS	Web User Interface
80	HTTPS	Redfish service root
443	Redfish	Redfish https with auth
623	RMCP+	IPMI
7582	KVM	Secure (SSL) KVM redirection
1900	UPNP	UPNP discovery
50000	UPNP	UPNP discovery
427	SLPD	Service Locator
123	NTP	Network Time Protocol
161	SNMP	SNMP incoming UDP requests
199	SNMP	SNMP incoming SMUX PDUs
546	DHCPv6	DHCPv6 messages
5124	CD Media redirection	CD media redirection secure (SSL) connections

6.4. Configuring a Static IP Address for the BMC

This section explains how to set a static IP address for the BMC. If your network does not support DHCP, you need to do this.

Use one of the methods described in the following sections:

- ▶ [Configuring a BMC Static IP Address Using ipmitool](#)
- ▶ [Configuring a BMC Static IP Address by Using the System BIOS](#)

6.4.1. Configuring a BMC Static Address by Using ipmitool

This section describes how to set a static IP address for the BMC from the Ubuntu command line.

 **Note**

If you cannot access the DGX B200 system remotely, connect a display (1440x900 or lower resolution) and keyboard directly to the DGX B200 system.

To view the current settings, enter the following command.

```
$ sudo ipmitool lan print 1
```

1. Set the IP address source to static.

```
$ sudo ipmitool lan set 1 ipsrc static
```

2. Set the appropriate address information.

- ▶ To set the IP address (Station IP address in the BIOS settings), enter the following and replace the italicized text with your information.

```
$ sudo ipmitool lan set 1 ipaddr <my-ip-address>
```

- ▶ To set the subnet mask, enter the following and replace the italicized text with your information.

```
$ sudo ipmitool lan set 1 netmask <my-netmask-address>
```

- ▶ To set the default gateway IP (Router IP address in the BIOS settings), enter the following and replace the italicized text with your information.

```
$ sudo ipmitool lan set 1 defgw ipaddr <my-default-gateway-ip-address>
```

6.4.2. Configuring a BMC Static IP Address by Using the System BIOS

This section describes how to set a static IP address for the BMC when you cannot access the DGX B200 System remotely, and this process involves setting the BMC IP address during system boot.

1. Connect a keyboard and display (1440 x 900 maximum resolution) to the DGX B200 system and turn on the DGX B200 system.
2. When you see the SBIOS version screen, press **Del** or **F2** to enter the **BIOS Setup Utility** screen.
3. On the **BIOS Setup Utility** screen, navigate to the **Server Mgmt** tab on the top menu. Scroll to **BMC network configuration** and press **Enter**.
4. Scroll to **Configuration Address Source** and press **Enter**. On the **Configuration Address Source** dialog, select **Static** and then press **Enter**.
5. Set the addresses for the Station IP address, Subnet mask, and Router IP address as needed by performing the following steps for each:
 1. Scroll to the specific item and press **Enter**.
 2. Enter the appropriate information in the dialog, and then press **Enter**.
6. When you finish making the changes, press **F10** to save and exit.

6.5. Changing the BMC Login Credentials

6.5.1. Username and Password Requirements

Refer to the following requirements for the username:

- ▶ Must be a string of 1 to 16 alphanumeric characters.
- ▶ Must start with an alphabetical character.
- ▶ Case-sensitive
- ▶ Special characters - (hyphen), _ (underscore), and @ (at sign) are allowed.


Refer to the following requirements for the password:

- ▶ Must be between 13 to 20 characters in length.
- ▶ Case-sensitive
- ▶ Special characters that must be preceded by a \ (backslash) character: ! "& ' () ; < > ` | } ~ \
- ▶ Special characters that do not require any special consideration: # \$ % * + , - . / : = ? @ [] ^ _ {

6.5.2. Procedure

To change your credentials or add or remove users, perform the following steps:

1. Select **Settings** from the left-side navigation menu.
2. Select the **User Management** card.



3. Click the help icon (?) for information about configuring users and creating passwords.
4. Log out and then log in with the new credentials.

6.6. Using the Remote Console


To use the remote console, perform the following steps:

1. Click **Remote Control** from the left-side navigation menu.
2. Click **Launch KVM** to start the remote KVM and access the DGX system console.


6.7. Setting Up Active Directory, LDAP, or E-Directory

To set up Active Directory, LDAP, or E-Directory, perform the following steps:

- From the side navigation menu, click **Settings** > **External User Services**.




- Click **Active Directory Settings** or **LDAP/E-Directory Settings** and follow the instructions.



6.8. Configuring Platform Event Filters

From the side navigation menu, click **Settings** and then click **Platform Event Filters**.




The Event Filters page shows all configured event filters and available slots. You can modify or add a new event filter entry on this page.

- To view available configured and unconfigured slots, click **All** in the upper-left corner of the page.
- To view available configured slots, click **Configured** in the upper-left corner of the page.
- To view available unconfigured slots, click **UnConfigured** in the upper-left corner of the page.
- To delete an event filter from the list, click the **x** icon.

6.9. Uploading or Generating SSL Certificates

You can set up a new certificate by generating a (self-signed) SSL or uploading an SSL (for example, using a Trusted CA-signed certificate).


From the side navigation menu, click **Settings > SSL Settings**.



Refer to the following sections for more information.

6.9.1. Viewing the SSL Certificate

To view the SSL certificate, on the SSL Setting page, click **View SSL Certificate**.



The View SSL Certificate page displays the following basic information about the uploaded SSL certificate:

- ▶ Certificate Version, Serial Number, Algorithm, and Public Key
- ▶ Issuer information
- ▶ Valid Date range
- ▶ Issued to information

6.9.2. Generating the SSL Certificate

Here is some information about generating an SSL certificate.

1. From the SSL Setting page, click **Generate SSL Certificate**.
2. Enter the information as described in the following table.

Table 3: SSL Certificate

Items	Description and Requirements
Common Name (CN)	The common name for which the certificate is to be generated. <ul style="list-style-type: none"> ▶ Maximum length of 64 alphanumeric characters. ▶ Special characters '#' and '\$' are not allowed.
Organization (O)	The name of the organization for which the certificate is generated. <ul style="list-style-type: none"> ▶ Maximum length of 64 alphanumeric characters. ▶ Special characters '#' and '\$' are not allowed.
Organization Unit (OU)	Overall organization section unit name for which the certificate is generated. <ul style="list-style-type: none"> ▶ Maximum length of 64 alphanumeric characters. ▶ Special characters '#' and '\$' are not allowed.
City or Locality (L)	City or Locality of the organization (mandatory) <ul style="list-style-type: none"> ▶ Maximum length of 64 alphanumeric characters. ▶ Special characters '#' and '\$' are not allowed.
State or Province (ST)	State or Province of the organization (mandatory) <ul style="list-style-type: none"> ▶ Maximum length of 64 alphanumeric characters. ▶ Special characters '#' and '\$' are not allowed.
Country (C)	Country code of the organization. <ul style="list-style-type: none"> ▶ Only two characters are allowed. ▶ Special characters are not allowed.
Email Address	Email address of the organization (mandatory)
Valid for	Enter a range from 1 to 3650 (days)
Key Length	Enter 4096.

3. To generate the new certificate, click **Save**.


6.9.3. Uploading the SSL Certificate

In BMC, you can upload your SSL certificate.

Make sure the certificate and key meet the following requirements:

- ▶ SSL certificates and keys must both use the .pem file extension.
- ▶ Private keys must not be encrypted.
- ▶ SSL certificates and keys must each be less than 3584 bits in size.
- ▶ SSL certificates must be current (not expired).

1. On the SSL Setting page, click **Upload SSL Certificate**.



2. Click the **New Certificate** folder icon, browse to locate the appropriate file, and select it.
3. Click the **New Private Key** folder icon, browse and locate the appropriate file, and select it.
4. Click **Save**.

6.9.4. Updating the SBIOS Certificate


The CA Certificate for the trusted CA that was used to sign the SSL certificate must be uploaded to allow the SBIOS to authenticate the certificate.

1. Obtain the CA certificate from the signing authority used to sign the SSL certificate.
2. Copy the CA certificate onto a USB thumb drive or to `/boot/efi` on the operating system.
3. Access a console from a locally connected keyboard and mouse or through the BMC remote console.
4. Reboot the server.
5. To enter BIOS setup menu, when prompted, press **DEL**.


Note

you may need to be logged in with admin privileges.


6. In the BIOS setup menu on the **Advanced** tab, select **Tls Auth Config**.



7. Select Server CA Configuration.




8. Select Enroll Cert.




9. Select Enroll Cert Using File.

10. Select the device where you stored the certificate.



11. Navigate the file structure and select the certificate.



Chapter 7. Managing Power Capping

The GPU has three sources of power limits:

- ▶ VBIOS: defines the maximum possible TGP (Total Graphics Power) value.
- ▶ The `nvidia-smi` tool: sets the power limit of the GPU through the host by users.
- ▶ SMBPBI: sets the power limit of the GPU via an out-of-band channel.

The GPU Performance Monitoring Unit (PMU) selects the most conservative policy to cap a system's power consumption.

7.1. Querying the Current GPU Power Limit

Use the following `curl` command to query the current GPU power limit:

```
curl -k -u <username>:<password> https://<bmc>/redfish/v1/Systems/HGX_<br>→Baseboard_0/Processors/GPU_SXM_<id>/EnvironmentMetrics
```

Where

- ▶ `<bmc>` is the BMC IP address.
- ▶ `<id>` is the GPU instance number of 1 to 8.

As shown in the following example output, the `Reading` field indicates the current power usage, and the `SetPoint` field indicates the current GPU power limit.

```
...  
"PowerLimitWatts": {  
    "AllowableMax": 1000,  
    "AllowableMin": 200,  
    "ControlMode": "Automatic",  
    "DefaultSetPoint": 700,  
    "Reading": 64.388,  
    "SetPoint": 700  
}  
...
```

7.2. Managing N+N Configuration (IPMI)

By default, a system will boot with three power supplies. To achieve the safe operation of an N+N configuration, you need to enable the power capping feature to limit the power consumed by the system.

Get the System Power Limit

```
ipmitool raw 0x3c 0x80 0x05
```

The format of the response is c8 32. To convert this value:

```
(0xc8 + 0x32 << 8) = 0x32c8 = 13000
```

If the feature is disabled, a value greater than 12,000 is returned.

Enable PSU Redundancy Support

To enable the PSU redundancy feature, set the power budget limit outside the actual system budget. The following example sets the power budget to 12 kW.

```
ipmitool raw x3c 0x81 0x05 0xE0 0x2E //Set 12 kW
```

Enable Power Capping Support

To operate the system lower than the maximum power budget the PSU can support, set a limit lower than:

```
ipmitool raw x3c 0x81 0x05 <MSB> <LSB>
```

The following example sets a limit of 6 KW (0x1770):

```
ipmitool raw 0x3c 0x81 0x05 0x70 0x17
```

Querying the GPU power limit using Redfish API shows 562 W:

```
...
"PowerLimitWatts": {
    "AllowableMax": 1000,
    "AllowableMin": 200,
    "ControlMode": "Automatic",
    "DefaultSetPoint": 700,
    "Reading": 64.335,
    "SetPoint": 562
}
...
```

7.3. Managing Power Capping Using Redfish API

To manage a system's maximum power consumption through power capping using Redfish API, refer to [Querying GPU Power Limit](#) and [Power Capping](#).

Chapter 8. Security

This section provides information about security measures in the NVIDIA DGX™ B200 system.

8.1. User Security Measures

The NVIDIA DGX B200 system is a specialized server designed to be deployed in a data center. It must be configured to protect the hardware from unauthorized access and unapproved use. The DGX B200 system is designed with a dedicated BMC Management Port and multiple Ethernet network ports.

When you install the DGX B200 system in the data center, follow the best practices established by your organization to protect against unauthorized access.

8.1.1. Securing the BMC Port

NVIDIA recommends connecting the BMC port in the DGX B200 system to a dedicated management network with firewall protection.

If remote access to the BMC is required, such as for a system hosted at a co-location provider, it should be accessed through a secure method that provides isolation from the internet, such as through a VPN server.

8.2. System Security Measures

This section provides information about the security measures incorporated in the NVIDIA DGX B200 system.

8.2.1. Secure Flash of DGX B200 Firmware

Secure Flash is implemented for the DGX B200 to prevent unsigned and unverified firmware images from being flashed onto the system.

8.2.2. Encryption

Here is some information about encrypting the DGX B200 firmware.

The firmware encryption algorithm is AES-CBC.

- ▶ The firmware encryption key strength is 128 bits or higher.
- ▶ Each firmware class uses a unique encryption key.
- ▶ Firmware decryption is performed by the same agent performing the signature check or a more trusted agent in the same COT.

8.2.3. NVIDIA System Manager Security

For information about security in NVIDIA System Management, refer to [NVSM documentation page](#).

8.3. Secure Data Deletion

This section explains how to securely delete data from the DGX B200 system SSDs to destroy all the stored data permanently.

This process performs a more secure SSD data deletion than merely deleting files or reformatting the SSDs.

8.3.1. Prerequisites

You need to prepare a bootable installation medium containing the current DGX OS Server ISO image.

Refer to [Reimaging the System](#) in the *NVIDIA DGX OS 7 User Guide* for information on the following topics:

- ▶ Obtaining the DGX OS ISO Image
- ▶ Booting the DGX OS ISO Image

8.3.2. Procedure

Here are the instructions to securely delete data from the DGX B200 system SSDs.

1. Boot the system from the ISO image, either remotely or from a bootable USB key.
2. At the GRUB menu, select:
 - ▶ (For DGX OS 7): **Rescue a broken system** and configure the locale and network information.
3. When prompted to select a root file system, select **Do not use a root file system** and then select **Execute a shell in the installer environment**.
4. Log in.
5. Run the following command to identify the devices available in the system:

```
nvme list
```

If the nvme-cli package is not installed, then install the CLI as follows and then run `nvme list`.

```
dpkg -i /usr/lib/live/mount/rootfs/filesystem.squashfs/curtin/repo/<nvme-  
cli-package.deb>
```

6. Perform a secure erase:

```
nvme format -s1 <device-path>
```

where <device-path> is the specific storage node listed in the previous step. For example, /dev/nvme0n1.

Chapter 9. Redfish APIs Support

The DGX System firmware supports Redfish APIs. Redfish is DMTF's standard set of APIs for managing and monitoring a platform. Redfish support is enabled by default in the DGX B200 BMC and the BIOS. Using the Redfish interface, administrator-privileged users can browse physical resources at the chassis and system level through the REST API interface. Redfish provides information that is categorized under a specific resource endpoint and Redfish clients can use the endpoints by using the following HTTP methods:

- ▶ GET
- ▶ POST
- ▶ PATCH
- ▶ PUT
- ▶ DELETE

Not all endpoints support all these operations. Refer to the Redfish JSON Schema for more information about the operations. The Redfish server follows the [DSP0266 1.7.0 Specification and Redfish Schema 2019.1](#) documentation. Redfish URLs are accessed using basic authentication and implementation so that IPMI users with required privileges can access the Redfish URLs.

9.1. Supported Redfish Features

Here is some information about the Redfish features supported in DGX B200.

The following features are supported:

- ▶ Manage user accounts, privileges, and roles
- ▶ Manager Sessions
- ▶ BMC configuration
- ▶ BIOS configuration
- ▶ BIOS boot order management
- ▶ Get PCIe device and functions inventory
- ▶ Get storage Inventory
- ▶ Get system component information and health (PSU, FAN, CPU, DIMM, and so on)
- ▶ Get sensor information (Thermal/Power/Cooling)
- ▶ BMC configuration change/BMC reset

- ▶ System/Chassis power operations
- ▶ Get health event log/advanced system event log
- ▶ Logging Service, which provides critical/informational severity events
- ▶ Event Services (SSE)
- ▶ Querying GPU power limit
- ▶ Power capping

Refer to the following documentation for more information:

- ▶ [DMTF Redfish specification](#)
- ▶ [DSP0266 1.7.0 specification](#)
- ▶ [Redfish Schema 2019.1 announcement from DMTF](#)

9.2. Connectivity Between the Host and BMC

You can configure internal network connectivity between the host and the BMC rather than using external network connectivity and routing traffic outside the host.

To configure internal network connectivity, you must configure an interface on the 169.254.0.0/255.255.0.0 network. The interface can then send and receive Redfish API traffic between the host and the BMC. The BMC is preconfigured to use the 169.254.0.17 IP address.

Run an `ifconfig` command like the following example to configure connectivity:

```
sudo ifconfig enx9638a3b292ec 169.254.0.18 netmask 255.255.0.0
```

Replace the network interface name and IP address in the preceding example according to your needs.

After you configure the network interface, you can use commands such as `curl` and `nvfwupd` with the 169.254.0.17 IP address to connect to the BMC and use the Redfish API.

The following example command shows the firmware versions:

```
nvfwupd -t ip=169.254.0.17 username=<bmc-user> password=<password> show_version
```

9.3. Redfish Examples

9.3.1. BMC Manager

- ▶ Accounts

The following `curl` command changes the password for the admin user.

```
curl -k -u <bmc-user>:<password> --request PATCH 'https://<bmc-ip-address>/redfish/v1/AccountService/Accounts/1' --header 'If-Match: *' --header 'Content-Type: application/json' --data-raw '{"Enabled": true, "Password": "DGXuser12345678!", "UserName": "admin", "RoleId": "Administrator", "Locked": false}'
```

► Reset BMC

The following curl command forces a reset of the DGX B200 BMC.

```
curl -k -u <bmc-user>:<password> --request POST --location 'https://<bmc-ip-address>/redfish/v1/Managers/BMC/Actions/Manager.Reset' --header
  'Content-Type: application/json' --data '{"ResetType": "ForceRestart"}'
```

► Reset BMC to factory defaults

The following curl command resets the BMC to factory defaults.

```
curl -k -u <bmc-user>:<password> --request POST --location 'https://<bmc-ip-address>/redfish/v1/Managers/BMC/Actions/Manager.ResetToDefaults' --
  --header 'Content-Type: application/json' --data '{"ResetType": "ResetAll"}'
```

9.3.2. Firmware Update

► Firmware inventory

```
curl -k -u <bmc-user>:<password> --request GET 'https://<bmc-ip-address>/redfish/v1/UpdateService/FirmwareInventory'
```

Example response:

```
{
  "@odata.context": "/redfish/v1/$metadata#SoftwareInventoryCollection.
  SoftwareInventoryCollection",
  "@odata.etag": "\"1683226281\"",
  "@odata.id": "/redfish/v1/UpdateService/FirmwareInventory",
  "@odata.type": "#SoftwareInventoryCollection.
  SoftwareInventoryCollection",
  "Description": "Collection of Firmware Inventory resources available
  to the UpdateService",
  "Members": [
    {
      "@odata.id": "/redfish/v1/UpdateService/FirmwareInventory/
  CPLDMB_0"
    },
    {
      "@odata.id": "/redfish/v1/UpdateService/FirmwareInventory/
  CPLDMID_0"
    },
    ...
  ],
  "Members@odata.count": 66,
  "Name": "Firmware Inventory Collection",
  "Oem": {
    "Ami": {
      "FirmwareInventory": [
        {
          "DataSourceUri": "/redfish/v1/UpdateService/
  FirmwareInventory"
        }
      ]
    }
  }
}
```

(continues on next page)

(continued from previous page)

```
  ↵FirmwareInventory/CPLDMB_0",
      "Name": "CPLDMB_0",
      "Version": "0.2.1.6"
    },
    {
      "DataSourceUri": "/redfish/v1/UpdateService/
 ↵FirmwareInventory/CPLDMID_0",
      "Name": "CPLDMID_0",
      "Version": "0.2.0.7"
    },
    // ...
  ]
}
}
```

► Update GPU tray components

To update the GPU tray components in your DGX B200 system, specify `HGX_0` as the target regardless of the GPU tray component that you want to update.

```
echo "{\"Targets\":[\"/redfish/v1/UpdateService/FirmwareInventory/HGX_0\"
→\"}]}" > parameters.json
curl -k -u <bmc-user>:<password> -H 'Expect:' --location --request POST
→https://<bmc-ip-address>/redfish/v1/UpdateService/upload -F
→'UpdateParameters=@parameters.json;type=application/json' -F
→'UpdateFile=@<fw_bundle>
```

For example, specify the `nvfw_DGX-HGX-B100-B200x8_0009_xxxxxxx.x.x_prod-signed.fwpx` firmware file.

► Update motherboard tray components

To update the motherboard tray components, you need to specify the component name as a target in a JSON file. The following example updates the host BMC:

```
echo "{\"Targets\":[\"/redfish/v1/UpdateService/FirmwareInventory/HostBMC_\"  
→\"0\"]}" > parameters.json  
curl -k -u <bmc-user>:<password> -H 'Expect:' --location --request POST  
→https://<bmc-ip-address>/redfish/v1/UpdateService/upload -F  
→'UpdateParameters=@parameters.json;type=application/json' -F  
→'UpdateFile=@<fw_bundle>'
```

The following targets are available:

- ▶ HostBMC_0 — This is the DGX B200 BMC.
- ▶ HostBIOS_0 — This is the DGX B200 BIOS.
- ▶ EROT_BMC_0 — This is the external root of trust for the host BMC.
- ▶ EROT_BIOS_0 — This is the external root of trust for the host BIOS.
- ▶ CPLDMID_0 — This is the midplane CPLD.
- ▶ CPLDMB_0 — This is the CPU tray CPLD.
- ▶ PSU_0 to PSU_5 — These are the PSUs.

- ▶ PCIeSwitch_0 and PCIeSwitch_1 — These are the Gen5 PCIe switches on the CPU tray.
- ▶ PCIeRetimer_0 and PCIeRetimer_1 — These are the PCIe retimers on the CPU tray.

To update a target, change the path /redfish/v1/UpdateService/FirmwareInventory/HostBMC_0 in the preceding example. For example, for CPU tray CPLD, specify /redfish/v1/UpdateService/FirmwareInventory/CPLDMB_0.

Make sure to specify the nvfw_DGX_0005_xxxxxx.x.x.fwpgk firmware file.

▶ Forced Update

The DGX B200 system component firmware is only updated if the incoming firmware version is newer than the existing version. To override this behavior and flash the component, specify the ForceUpdate field and set it to true.

```
curl -k -u <bmc-user>:<password> --request PATCH 'https://<bmc-ip-address>/redfish/v1/UpdateService' --header 'If-Match: *' --header 'Content-Type: application/json' --data-raw '{"HttpPushUriOptions": {"ForceUpdate": true}}'
```

On success, the command returns a 204 HTTP status code. If you attempt to set the flag to the currently set value, the command returns a 400 HTTP status code.

To get the value of the ForceUpdate parameter:

```
curl -k -u <bmc-user>:<password> --request GET 'https://<bmc-ip-address>/redfish/v1/UpdateService'
```

▶ Firmware Update Activation

To activate the firmware update, refer to [Firmware Update Activation](#) in the *NVIDIA DGX B200 Firmware Update Guide* for more information.

▶ Background Copy

After the firmware image is successfully booted, use the background copy API to synchronize firmware images between partitions.

```
curl -k -u <bmc-user>:<password> --request POST --location 'https://<bmc-ip-address>/redfish/v1/UpdateService/Actions/Oem/NvidiaUpdateService.CommitImage' \
--header 'Content-Type: application/json' \
--data '{
    "Targets": ["/redfish/v1/UpdateService/FirmwareInventory/HostBIOS_0"]
}'
```

9.3.3. BIOS Settings

- ▶ Supported BIOS attributes
- ▶ Get a list of all the attributes that your BIOS supports:

```
curl -k -u <bmc-user>:<password> --location --request GET 'https://<bmc-ip-address>/redfish/v1/Registries'
```

One of the Registries in the list is your BIOS attribute registry. The format is BiosAttributeRegistry<version><version>. For example, for BIOS 0.1.6, the registry is BiosAttributeRegistry106.1.0.6.

- ▶ Get the URI of the BIOS registry:

```
curl -k -u <bmc-user>:<password> --location --request GET 'https://
↪<bmc-ip-address>/redfish/v1/Registries/BiosAttributeRegistry016.0.1.
↪6/'
```

The response includes the location of the JSON file that describes all the BIOS attributes. Under Location, the Uri is specified. For example, Uri": "/redfish/v1/Registries/BiosAttributeRegistry106.1.0.6.

- ▶ Get the JSON file with the registry of all your BIOS attributes:

```
curl -k -u <bmc-user>:<password> --location --request GET 'https://
↪<bmc-ip-address>/redfish/v1/Registries/BiosAttributeRegistry106.en-
↪US.1.0.6.json' --output BiosAttributeRegistry106.en-US.1.0.6.json
```

Each attribute name has a default value, display name, help text, a read-only indicator, and an indicator of whether a reset is required to take effect.

- ▶ To get the current BIOS settings:

```
curl -k -u <bmc-user>:<password> --location --request GET 'https://<bmc-
↪ip-address>/redfish/v1/Systems/DGX/Bios'
```

Match the attribute name with the value in the registry for a description.

Example response:

```
"Description": "Current BIOS Settings",
"Id": "Bios",
"Name": "Current BIOS Settings"
...
```

- ▶ To change an attribute in the future BIOS settings, PATCH the SD URI and specify the attribute name with the new value. You can change more than one attribute at a time.

For example, the following PATCH request specifies how the system responds when the SEL log is full:

```
curl -k -u <bmc-user>:<password> --location --request PATCH 'https://<bmc-
↪ip-address>/redfish/v1/Systems/DGX/Bios/SD' -H 'Content-Type:
↪application/json' -H 'If-Match:' --data-raw '{"Attributes": {"IPMI002
↪": "IPMI002DoNothing", "IPMI201": "IPMI201DoNotLoganymore"}}'
```

Example response:

```
"Description": "Future BIOS Settings",
"Id": "SD",
"Name": "Future BIOS Settings"
...
```

Note

All attribute changes to the BIOS require a power cycle to take effect. When changing the attributes is followed by a BIOS update, an additional power cycle is needed to apply the changes.

9.3.4. Modifying the Boot Order Using Redfish

To modify the boot order on DGX B200 systems using Redfish APIs, follow the steps described in this procedure.

1. Read the current boot order.

From any system in the same network as the BMC, run the following `curl` command to get the current boot order:

```
$ curl -k -u <BMC username>:<BMC password> https://<BMC_IP_address>/
  ↵redfish/v1/Systems/DGX/SD -H "content-type:application/json" -X GET -s
  ↵| jq .Boot.BootOrder
```

```
[  
  "Boot0000",  
  "Boot000F",  
  "Boot0004",  
  "Boot0005",  
  "Boot0006",  
  "Boot0007",  
  "Boot0008",  
  "Boot0009",  
  "Boot000A",  
  "Boot0010"  
]
```

2. Identify the available boot devices.

To show more information about the boot devices in step 1, such as `Boot0000`, `Boot000F`, and `Boot0004`, run the following command:

```
$ curl -k -u <BMC username>:<BMC password> https://<BMC_IP_address>/
  ↵redfish/v1/Systems/DGX/BootOptions/00{0,1}{0,4,5,6,7,8,9,A,F} -H
  ↵"content-type:application/json" -X GET -s | jq |grep -e
  ↵"UefiDevicePath\|Name"
```

```
"@odata.etag": "\"1696896625\"",  
"DisplayName": "DGX OS",  
"Name": "Boot0000",  
"UefiDevicePath": "HD(1,GPT,159C2E52-2329-40AC-9103-6C28DC1528B8,0x800,  
  ↵0x100000)/\EFI\UBUNTU\SHIMX64.EFI"  
"@odata.etag": "\"1696896625\"",  
"DisplayName": "UEFI: PXE IPv4 Intel(R) Ethernet Controller X550",  
"Name": "Boot0004",  
"UefiDevicePath": "PciRoot(0x0)/Pci(0x10,0x0)/Pci(0x0,0x0)/  
  ↵MAC(5CFF35FBDA09,0x1)/IPv4(0.0.0.0,0x0,DHCP,0.0.0.0,0.0.0.0,0.0.0.0)"
```

(continues on next page)

(continued from previous page)

```

"@odata.etag": "\"1696896625\"",
"DisplayName": "UEFI: PXE IPv4 Nvidia Network Adapter - B8:3F:D2:E7:B1:6C
",
"Name": "Boot0005",
"UefiDevicePath": "PciRoot(0x20)/Pci(0x1,0x0)/Pci(0x0,0x0)/Pci(0x0,0x0)/
Pci(0x0,0x0)/Pci(0x0,0x0)/Pci(0x0,0x0)/MAC(B83FD2E7B16C,0x1)/IPv4(0.0.0.
0,0x0,DHCP,0.0.0.0,0.0.0.0,0.0.0.0)"
"@odata.etag": "\"1696896625\"",
"DisplayName": "UEFI: PXE IPv4 Nvidia Network Adapter - B8:3F:D2:E7:B1:6D
",
"Name": "Boot0006",
"UefiDevicePath": "PciRoot(0x20)/Pci(0x1,0x0)/Pci(0x0,0x0)/Pci(0x0,0x0)/
Pci(0x0,0x0)/Pci(0x0,0x0)/Pci(0x0,0x1)/MAC(B83FD2E7B16D,0x1)/IPv4(0.0.0.
0,0x0,DHCP,0.0.0.0,0.0.0.0,0.0.0.0)"
"@odata.etag": "\"1696896625\"",
"DisplayName": "UEFI: PXE IPv4 Nvidia Network Adapter - B8:3F:D2:E7:B0:9C
",
"Name": "Boot0007",
"UefiDevicePath": "PciRoot(0x120)/Pci(0x1,0x0)/Pci(0x0,0x0)/Pci(0x0,0x0)/
Pci(0x0,0x0)/Pci(0x0,0x0)/Pci(0x0,0x0)/MAC(B83FD2E7B09C,0x1)/IPv4(0.0.0.
0,0x0,DHCP,0.0.0.0,0.0.0.0,0.0.0.0)"
"@odata.etag": "\"1696896625\"",
"DisplayName": "UEFI: PXE IPv4 Nvidia Network Adapter - B8:3F:D2:E7:B0:9D
",
"Name": "Boot0008",
"UefiDevicePath": "PciRoot(0x120)/Pci(0x1,0x0)/Pci(0x0,0x0)/Pci(0x0,0x0)/
Pci(0x0,0x0)/Pci(0x0,0x0)/Pci(0x0,0x1)/MAC(B83FD2E7B09D,0x1)/IPv4(0.0.0.
0,0x0,DHCP,0.0.0.0,0.0.0.0,0.0.0.0)"
"@odata.etag": "\"1696896625\"",
"DisplayName": "UEFI: PXE IPv4 Intel(R) Ethernet Network Adapter E810-C-Q2
",
"Name": "Boot0009",
"UefiDevicePath": "PciRoot(0x160)/Pci(0x5,0x0)/Pci(0x0,0x0)/
MAC(6CFE543D8F48,0x1)/IPv4(0.0.0.0,0x0,DHCP,0.0.0.0,0.0.0.0,0.0.0.0)"
"@odata.etag": "\"1696896625\"",
"DisplayName": "UEFI: PXE IPv4 Intel(R) Ethernet Network Adapter E810-C-Q2
",
"Name": "Boot000A",
"UefiDevicePath": "PciRoot(0x160)/Pci(0x5,0x0)/Pci(0x0,0x1)/
MAC(6CFE543D8F49,0x1)/IPv4(0.0.0.0,0x0,DHCP,0.0.0.0,0.0.0.0,0.0.0.0)"
"@odata.etag": "\"1696896625\"",
"DisplayName": "ubuntu",
"Name": "Boot000F",
"UefiDevicePath": "HD(1,GPT,1E0EFF2A-2BF3-4DC6-8757-4075B1E5343D,0x800,
0x100000)/\EFI\UBUNTU\SHIMX64.EFI"
"@odata.etag": "\"1696896625\"",
"DisplayName": "UEFI: PXE IPv4 American Megatrends Inc.",
"Name": "Boot0010",
"UefiDevicePath": "PciRoot(0x0)/Pci(0x14,0x0)/USB(0xA,0x0)/USB(0x2,0x1)/
MAC(4E2A712C2451,0x0)/IPv4(0.0.0.0,0x0,DHCP,0.0.0.0,0.0.0.0,0.0.0.0)"

```

Where

- ▶ The DisplayName string is the name of the drive or network adapter.
- ▶ The Name string is the boot device name.
- ▶ The MAC (<address>, 0x1) value for the UefiDevicePath string is the corresponding MAC address.
- ▶ The @odata.etag string is the etag number.

Identify the following information from the JSON output for the next step:

- ▶ The name of the device to be the boot device.
- ▶ The etag number to compose the header.

3. Update the boot order.

The following command uses the PATCH method to modify the BootOrder settings, specifying the etag number and boot device names from step 2. The command generates a new order list for BootOrder, which affects the next boot of the system.

```
$ curl -k -u <BMC username>:<BMC password> https://<BMC_IP_address>/redfish/v1/Systems/DGX/SD -H "content-type:application/json" -H 'if-None-Match: "@odata.etag": "1697483651"' --data '{"Boot":{"BootOrder": ["Boot0004", "Boot0000", "Boot0005", "Boot0006", "Boot0007", "Boot0008", "Boot0009", "Boot000A", "Boot000F", "Boot0010"]}}' -X PATCH
```


4. Confirm the boot order.

Repeat the command in step 1 to ensure the BootOrder settings are as expected. Note that the Boot0004 boot device is now at the top and the system will boot from the on-board RJ-45 network interface.

```
$ curl -k -u <BMC username>:<BMC password> https://<BMC_IP_address>/redfish/v1/Systems/DGX/SD -H "content-type:application/json" -X GET -s | jq .Boot.BootOrder
```

```
[  
  "Boot0004",  
  "Boot0000",  
  "Boot0005",  
  "Boot0006",  
  "Boot0007",  
  "Boot0008",  
  "Boot0009",  
  "Boot000A",  
  "Boot000F",  
  "Boot0010"]
```

Upon reboot, the system should attempt to boot from the network using the correct network interface:



This boot order change will remain until the next boot order update, which can be done by resetting the SBIOS or running this procedure again.

9.3.5. Telemetry

- ▶ GPU tray sensors

```
curl -k -u <bmc-user>:<password> --location --request GET 'https://<bmc-
<ip-address>/redfish/v1/TelemetryService/MetricReportDefinitions/HGX_
PlatformEnvironmentMetrics_0'
```

- ▶ DGX platform sensors

```
curl -k -u <bmc-user>:<password> --location --request GET 'https://<bmc-
<ip-address>/redfish/v1/Chassis/DGX/Sensors'
```

The endpoint returns 75 members at a time. To page through the results, use the URI in the `Members@odata.nextLink` field. For example, `/redfish/v1/Chassis/DGX/Sensors?skip=75`.

9.3.6. Chassis

- ▶ Chassis Restart (IPMI chassis power cycle)

```
curl -k -u <bmc-user>:<password> --request POST --location 'https://<bmc-
<ip-address>/redfish/v1/Systems/DGX/Actions/ComputerSystem.Reset' --
header 'Content-Type: application/json' --data '{"ResetType": "ForceRestart"}'
```

- ▶ Chassis Start (IPMI chassis power on)

```
curl -k -u <bmc-user>:<password> --request POST --location 'https://<bmc-
<ip-address>/redfish/v1/Systems/DGX/Actions/ComputerSystem.Reset' --
header 'Content-Type: application/json' --data '{"ResetType": "On"}'
```

- ▶ Chassis Graceful Restart (IPMI chassis soft off, IPMI chassis power on)

```
curl -k -u <bmc-user>:<password> --request POST --location 'https://<bmc-
<ip-address>/redfish/v1/Systems/DGX/Actions/ComputerSystem.Reset' --
header 'Content-Type: application/json' --data '{"ResetType": "GracefulRestart"}'
```

- ▶ Chassis Off (IPMI chassis power off)

```
curl -k -u <bmc-user>:<password> --request POST --location 'https://<bmc-ip-address>/redfish/v1/Systems/DGX/Actions/ComputerSystem.Reset' --header 'Content-Type: application/json' --data '{"ResetType": "ForceOff"}'
```

- ▶ Chassis Off Gracefully (IPMI chassis soft off)

```
curl -k -u <bmc-user>:<password> --request POST --location 'https://<bmc-ip-address>/redfish/v1/Systems/DGX/Actions/ComputerSystem.Reset' --header 'Content-Type: application/json' --data '{"ResetType": "GracefulShutdown"}'
```

- ▶ Chassis Power Cycle (IPMI chassis power off, IPMI chassis power on)

```
curl -k -u <bmc-user>:<password> --request POST --location 'https://<bmc-ip-address>/redfish/v1/Systems/DGX/Actions/ComputerSystem.Reset' --header 'Content-Type: application/json' --data '{"ResetType": "PowerCycle"}'
```

Note

The ForceRestart, GracefulRestart, and GracefulShutdown reset actions on HMC are not supported for security reasons.

9.3.7. SEL Logs

To view all the SEL entries using Redfish:

```
curl -k -u <bmc-user>:<password> --location --request GET 'https://<bmc-ip-address>/redfish/v1/Managers/BMC/LogServices/SEL/Entries'
```

The endpoint returns 75 members at a time. To page through the results, use the URI in the `Members@odata.nextLink` field. For example, `/redfish/v1/Managers/BMC/LogServices/SEL/Entries?$skip=75`.

9.3.8. Virtual Image

1. Enable remote virtual media and verify it is active.

```
curl -s -k -u <bmc-user>:<password> -H content-type:application/json -X POST https://<bmc-ip-address>/redfish/v1/Managers/BMC/Actions/Oem/AMIVirtualMedia.EnableRMedia -d '{"RMediaState": "Enable"}' -H 'If-None-Match: not_match' | jq
```

Example response:

```
{
  "error": {
    "@Message.ExtendedInfo": [
```

(continues on next page)

(continued from previous page)

```
{
    "@odata.type": "#Message.v1_0_8.Message",
    "Message": "EnableRMedia action has been initiated successfully.
    ↪Please allow upto 4-5 secs and verify the value of RMediaStatus OEM
    ↪property in /redfish/v1/Managers/BMC instance",
    "MessageArgs": [
        "EnableRMedia",
        "RMediaStatus OEM",
        "/redfish/v1/Managers/BMC"
    ],
    "MessageId": "Ami.1.0.DelayInActionCompletion",
    "Resolution": "Check the property value update after 4-5 seconds",
    "Severity": "OK"
}
],
"code": "Ami.1.0.DelayInActionCompletion",
"message": "EnableRMedia action has been initiated successfully.
    ↪Please allow upto 4-5 secs and verify the value of RMediaStatus OEM
    ↪property in /redfish/v1/Managers/BMC instance"
}
}
```

```
curl -k -u <bmc-user>:<password> -H "Content-Type: application/json"
    ↪https://<bmc-ip-address>/redfish/v1/Managers/BMC/VirtualMedia | jq
```

Example response:

% Total	% Received	% Xferd	Average Speed	Time	Time	Time
100	616	100	616	0	0	925
↪ 926				Dload	Upload	Total
						Spent
						Left
						Speed

```
{
    "@odata.context": "/redfish/v1/$metadata#VirtualMediaCollection.
    ↪VirtualMediaCollection",
    "@odata.etag": "\"1749834214\"",
    "@odata.id": "/redfish/v1/Managers/BMC/VirtualMedia",
    "@odata.type": "#VirtualMediaCollection.VirtualMediaCollection",
    "Description": "Collection of Virtual Media redirected to host via this
    ↪Manager",
    "Members": [
        {
            "@odata.id": "/redfish/v1/Managers/BMC/VirtualMedia/CD1"
        },
        {
            "@odata.id": "/redfish/v1/Managers/BMC/VirtualMedia/CD2"
        },
        {
            "@odata.id": "/redfish/v1/Managers/BMC/VirtualMedia/CD3"
        },
        {
            "@odata.id": "/redfish/v1/Managers/BMC/VirtualMedia/CD4"
        }
    ]
}
```

(continues on next page)

(continued from previous page)

```

        },
        ],
        "Members@odata.count": 4,
        "Name": "Virtual Media Collection"
    }
}
```

2. To mount the ISO file over NFS, send the Redfish POST request.

```
curl -k -u <bmc-user>:<password> -H "Content-Type: application/json" -  
request POST https://<bmc-ip-address>/redfish/v1/Managers/BMC/  
→VirtualMedia/CD1/Actions/VirtualMedia.InsertMedia -d '{"Image": "//  
→<server-ip-address>/<ISO-image-path>/<ISO-image-filename>",  
→"TransferProtocolType": "NFS"}' | jq
```

Example response:

```
% Total      % Received % Xferd  Average Speed   Time      Time      Time  
→Current  
          Dload  Upload   Total   Spent   Left   Speed  
100  344  100  231  100  113    236    115 --:--:-- --:--:-- --:--:  
→ 352  
{  
    "@odata.context": "/redfish/v1/$metadata#Task.Task",  
    "@odata.id": "/redfish/v1/TaskService/Tasks/1",  
    "@odata.type": "#Task.v1_4_2.Task",  
    "Description": "Task for InsertMedia Action",  
    "Id": "1",  
    "Name": "InsertMedia Action",  
    "TaskState": "New"  
}
```

3. Verify the task is complete and then check that the image is mounted.

```
curl -k -u <bmc-user>:<password> -H "Content-Type: application/json"  
→https://<bmc-ip-address>/redfish/v1/Managers/BMC/VirtualMedia/CD1 | jq
```

Example response:

```
% Total      % Received % Xferd  Average Speed   Time      Time      Time  
→Current  
          Dload  Upload   Total   Spent   Left   Speed  
100  1115  100  1115     0      0    1413      0 --:--:-- --:--:-- --:--:  
→1414  
{  
    "@odata.context": "/redfish/v1/$metadata#VirtualMedia.VirtualMedia",  
    "@odata.etag": "\"1749843453\"",  
    "@odata.id": "/redfish/v1/Managers/BMC/VirtualMedia/CD1",  
    "@odata.type": "#VirtualMedia.v1_3_2.VirtualMedia",  
    "Actions": {  
        "#VirtualMedia.EjectMedia": {  
            "@Redfish.ActionInfo": "/redfish/v1/Managers/BMC/VirtualMedia/CD1/  
→EjectMediaActionInfo",  
            "target": "/redfish/v1/Managers/BMC/VirtualMedia/CD1/Actions/
```

(continues on next page)

(continued from previous page)

```

    ↵"VirtualMedia.EjectMedia"
    },
    "#VirtualMedia.InsertMedia": {
        "@Redfish.ActionInfo": "/redfish/v1/Managers/BMC/VirtualMedia/CD1/
    ↵InsertMediaActionInfo",
        "target": "/redfish/v1/Managers/BMC/VirtualMedia/CD1/Actions/
    ↵VirtualMedia.InsertMedia"
    }
},
"ConnectedVia": "URI",
"Description": "Virtual Media Instance redirected to host via this
    ↵Manager",
"Id": "CD1",
"Image": "//server-ip-address:0/home/nvidia/images/DGXOS-7.0.0-2024-08-
    ↵13-07-50-56.iso",
"ImageName": "DGXOS-7.0.0-2024-08-13-07-50-56.iso",
"Inserted": true,
"MediaTypes": [
    "CD"
],
"Name": "CD1",
"Oem": {
    "Ami": {
        "@odata.type": "#AMIVirtualMedia.v1_0_0.AMIVirtualMedia",
        "RedirectionStatus": "Redirection Started"
    }
},
"TransferMethod": "Stream",
"TransferProtocolType": "NFS",
"UserName": "",
"WriteProtected": true
}
}

```

4. From Redfish, you can configure the SBIOS boot order and set ''UEFI: AMI Virtual CDROM0 1.00'' to the boot device, following the steps in [Modifying the Boot Order Using Redfish](#).

9.3.9. Backing Up and Restoring BMC Configurations

In addition to using the Web UI to back up and restore the BMC configuration, you can use Redfish API with the following approach:

1. Install a security AES key in the BMC.
2. Back up the BMC configuration.
3. Restore the BMC configuration using a backup file.

The BMC automatically reboots when you perform a configuration restore.

9.3.9.1 Backing Up the BMC Configuration

1. Generate an AES key and save it to a .bin file.

```
openssl rand -out aes_key.bin 32
```

2. Upload the AES key.

```
curl -s -k -u <username>:<password> --location --request POST 'https://<bmcip>/redfish/v1/Managers/BMC/Actions/Oem/NvidiaManager.UploadAESKey' --form 'AESKey=@aes_key.bin' | jq
```

A successful command returns a 204 HTTP status code.

3. Back up the BMC configuration by creating a backup file, for example, bmc-config.bak.

```
curl -s -k -u <username>:<password> POST 'https://<bmcip>/redfish/v1/Managers/BMC/Actions/Oem/NvidiaManager.BackupConfig' -H 'Content-Type: application/json' --data-raw '{ "BackupFeatures" : [ "NTP", "Network and Services", "Syslog", "Authentication", "SNMP", "IPMI", "KVM" ] }' > bmc-config.bak
```

9.3.9.2 Restoring the BMC Configuration

Note

You must perform a factory reset to restore the default settings before restoring the BMC configuration.

1. Upload the AES key generated previously.

```
curl -s -k -u <username>:<password> --location --request POST 'https://<bmcip>/redfish/v1/Managers/BMC/Actions/Oem/NvidiaManager.UploadAESKey' --form 'AESKey=@aes_key.bin' | jq
```

A successful command returns a 204 HTTP status code.

2. Restore the BMC configuration from a backup file, for example, bmc-config.bak.

```
curl -s -k -u <username>:<password> --location --request POST 'https://<bmcip>/redfish/v1/Managers/BMC/Actions/Oem/NvidiaManager.RestoreConfig' --form 'conf_file=@"bmc-config.bak"' | jq
```

9.3.10. Collecting BMC Debug Data

1. Create a request for BMC to start collecting debug data:

```
curl -k -u <bmc-user>:<password> --request POST --location 'https://<bmc-ip-address>/redfish/v1/Managers/BMC/LogServices/DiagnosticLog/Actions/LogService.CollectDiagnosticData' -H 'Content-Type: application/json' --data-raw '{ "DiagnosticDataType" : "OEM", "OEMDiagnosticDataType": "ALL" }' | jq
```

Example response:

```
{
  "@odata.context": "/redfish/v1/$metadata#Task.Task",
  "@odata.id": "/redfish/v1/TaskService/Tasks/2",
  "@odata.type": "#Task.v1_4_2.Task",
  "Description": "Task for Manager CollectDiagnosticData",
  "Id": "2",
  "Name": "Manager CollectDiagnosticData",
  "TaskState": "New"
}
```

2. Change the task number to the appropriate task Id returned from step 1, and monitor the task for completion until PercentComplete reaches 100.

```
curl -k -u <bmc-user>:<password> --request GET 'https://<bmc-ip-address>/
  ↵redfish/v1/TaskService/Tasks/2' | jq
```

Example response:

```
{
  "@odata.context": "/redfish/v1/$metadata#Task.Task",
  "@odata.etag": "\"1723565599\"",
  "@odata.id": "/redfish/v1/TaskService/Tasks/2",
  "@odata.type": "#Task.v1_4_2.Task",
  "Description": "Task for Manager CollectDiagnosticData",
  "EndTime": "2024-08-13T16:28:15+00:00",
  "Id": "2",
  "Messages": [
    {
      "@odata.type": "#Message.v1_0_8.Message",
      "Message": "Indicates that a DiagnosticDump of was created at /
        ↵redfish/v1/Managers/BMC/LogServices/DiagnosticLog/Attachment/nvidiadiag-
        ↵HT9buy.tar.gz",
      "MessageArgs": [
        "/redfish/v1/Managers/BMC/LogServices/DiagnosticLog/
        ↵Attachment/nvidiadiag-HT9buy.tar.gz"
      ],
      "MessageId": "Ami.1.0.0.DiagnosticDumpCreated",
      "Resolution": "None",
      "Severity": "Warning"
    },
    {
      "@odata.type": "#Message.v1_0_8.Message",
      "Message": "Task /redfish/v1/Managers/BMC/LogServices/
        ↵DiagnosticLog/Actions/LogService.CollectDiagnosticData has completed.",
      "MessageArgs": [
        "/redfish/v1/Managers/BMC/LogServices/DiagnosticLog/Actions/
        ↵LogService.CollectDiagnosticData"
      ],
      "MessageId": "Task.1.0.Completed",
      "Resolution": "None",
      "Severity": "OK"
    }
  ],
  "PercentComplete": 100
}
```

(continues on next page)

(continued from previous page)

```

    "Name": "Manager CollectDiagnosticData",
    "PercentComplete": 100,
    "StartTime": "2024-08-13T16:13:20+00:00",
    "TaskState": "Completed",
    "TaskStatus": "OK"
}

```

3. After the TaskState field indicates Completed, use the path provided by MessageArgs to download the attachment.

In the following command, substitute nvidiadiag-HT9buy.tar.gz (an example filename) with the actual attachment generated in step 2.

```
curl -k -u <bmc-user>:<password> --request GET 'https://<bmc-ip-address>/
˓→redfish/v1/Managers/BMC/LogServices/DiagnosticLog/Attachment/nvidiadiag-
˓→HT9buy.tar.gz' --output nvidiadiag-HT9buy.tar.gz
```

9.3.11. Clear BIOS and Reset to Factory Defaults

To clear the BIOS and reset the system to factory defaults:

```
curl -k -u <username>:<password> --request POST --location 'https://<bmcip>/
˓→redfish/v1/UpdateService/Actions/Oem/NvidiaUpdateService.ClearNVRAM' --
˓→header 'Content-Type: application/json' \
--data '{"Targets": ["/redfish/v1/UpdateService/FirmwareInventory/HostBIOS_0
˓→"]}'
```

9.3.12. Querying GPU Power Limit

- To query the current GPU power limit:

```
curl -k -u <username>:<password> https://<bmc>/redfish/v1/Systems/HGX_0/
˓→Baseboard_0/Processors/GPU_SXM_<id>/EnvironmentMetrics
```

Where

- <bmc> is the BMC IP address.
- <id> is the GPU instance number of 1 to 8.

As shown in the following example output, the Reading field indicates the current power usage, and the SetPoint field indicates the current GPU power limit.

```

...
"PowerLimitWatts": {
    "AllowableMax": 1000,
    "AllowableMin": 200,
    "ControlMode": "Automatic",
    "DefaultSetPoint": 700,
    "Reading": 64.388,
    "SetPoint": 700
}

```

(continues on next page)

(continued from previous page)

}

...

9.3.13. Power Capping

9.3.13.1 Services

To discover the available services:

```
curl -k -u <bmc-user>:<password> https://<bmcip>/redfish/v1/Managers/BMC/
  ↪NodeManager
```

Example response:

```
{
  "@odata.context": "/redfish/v1/$Metadata#NvidiaNodeManager.
  ↪NvidiaNodeManager",
  "@odata.id": "/redfish/v1/Managers/BMC/NodeManager",
  "@odata.type": "#NvidiaNodeManager.v1_0_0.NvidiaNodeManager",
  "ActiveDomain": {
    "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/12"
  },
  "Id": "NvidiaNodeManager",
  "Name": "NvidiaNodeManager",
  "NvidiaNmDomains": {
    "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains"
  },
  "NvidiaPSUPolicies": {
    "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/PSUPolicies"
  }
}
```

9.3.13.2 Domains

There are several predefined domains. If no domains are set, the default domains are shown.

- To get a list of domains:

The domain list contains a set of fixed domains (for example, 0-5) followed by custom domains starting at 10.

Request:

```
curl -k -u <bmc-user>:<password> https://<bmcip>/redfish/v1/Managers/BMC/
  ↪NodeManager/Domains
```

Example response:

```
{
  "@odata.context": "/redfish/v1/$Metadata#NvidiaNmDomainCollection.
  ↪NvidiaNmDomainCollection",
  "@odata.id": "/redfish/v1/Managers/BMC/NvidiaNmDomainCollection",
  "@odata.type": "#NvidiaNmDomainCollection.NvidiaNmDomainCollection",
```

(continues on next page)

(continued from previous page)

```

"Members": [
    {
        "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/0"
    },
    {
        "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/1"
    },
    {
        "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/4"
    },
    {
        "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/2"
    },
    {
        "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/3"
    },
    {
        "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/5"
    },
    {
        "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/10"
    },
    {
        "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/11"
    },
    {
        "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/12"
    },
    {
        "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/13"
    }
],
"Members@odata.count": 10,
"Name": "NvidiaNmDomainCollection"
}

```

- ▶ To view domain policies:

A domain policy describes the limits for the domain and the associated policies

The most important fields are:

- ▶ **Min and Max:** Power Limits, defined under **Capabilities**. The values reflected are system-specific.
- ▶ **Status:** Indicates if domain policy is enabled.
- ▶ **Policy State:** Indicates whether this policy is currently in effect.

```
curl -k -u <bmc-user>:<password> https://<bmcip>/redfish/v1/Managers/BMC/
    NodeManager/Domains/<DomainID>
```

For example, to view policies in domain 0:

```
curl -k -u <bmc-user>:<password> https://<bmcip>/redfish/v1/Managers/BMC/
  ↳NodeManager/Domains/0
```

Example response:

```
{
  "@odata.context": "/redfish/v1/$Metadata#NvidiaNmDomain.NvidiaNmDomain",
  "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/0",
  "@odata.type": "#NvidiaNmDomain.v1_4_0.NvidiaNmDomain",
  "Capabilities": {
    "Max": 16500,
    "MaxCorrectionTimeInMs": 2000,
    "MaxStatisticsReportingPeriod": "2000",
    "Min": 5000,
    "MinCorrectionTimeInMs": 1000,
    "MinStatisticsReportingPeriod": "1000"
  },
  "Id": "0",
  "Name": "protection",
  "Policies": {
    "@odata.context": "/redfish/v1/$Metadata#NvidiaNmPolicyCollection.
  ↳NvidiaNmPolicyCollection",
    "@odata.type": "#NvidiaNmPolicyCollection.NvidiaNmPolicyCollection"
  },
  "Members": [
    {
      "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/
  ↳Domains/0/Policies/0"
    },
    {
      "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/
  ↳Domains/0/Policies/1"
    },
    {
      "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/
  ↳Domains/0/Policies/2"
    }
  ],
  "Name": "NvidiaNmPolicyCollection"
},
"Status": {
  "State": "Enabled"
}
}
```

- To view a policy within a domain:

Each domain has a set of policies that define how to manage each component. Power is divided based on a percentage with a component not allowed to exceed a specific budget.

Currently, three policies are possible. The number of policies varies for different systems depending on whether the entity supports power management.

The most important fields are:

- ▶ **Limit**: Indicates the max power that can be allocated to this device.
- ▶ **PercentageofDomainBudget**: How much of the budget can be allocated.
- ▶ **Status**: Whether the policy is in effect. This determined by the Domain Policy State

In general, the algorithm always uses **PercentageofDomainBudget**.

```
curl -k -u <bmc-user>:<password> https://<bmcip>/redfish/v1/Managers/BMC/
  ↳NodeManager/Domains/0/Policies/<PolicyID>
```

For example, to view policy 0 in domain 0:

```
curl -k -u <bmc-user>:<password> https://<bmcip>/redfish/v1/Managers/BMC/
  ↳NodeManager/Domains/0/Policies/0
```

Example response:

```
{
  "@odata.context": "/redfish/v1/$Metadata#NvidiaNmPolicy.NvidiaNmPolicy",
  "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/0/Policies/0",
  "@odata.type": "#NvidiaNmPolicy.v1_2_0.NvidiaNmPolicy",
  "AssociatedDomainID": {
    "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/0"
  },
  "ComponentId": "COMP_MEMORY",
  "Id": "0",
  "Limit": 800,
  "Name": "0",
  "PercentageOfDomainBudget": 15,
  "Status": {
    "State": "Disabled"
  }
}
```

In this example, policy 0 defines the percentage of the budget for domain 0. The CPU budget for both sockets is 800 W, which is equally divided. The **PercentageOfDomainBudget** field, which indicates how much of the overall budget will be allocated to the CPUs, shows 15 percent for this example.

9.3.13.3 Custom Policies

To add a custom policy, use the following template and specify values for the highlighted fields. Custom domain ID starts from 10.

The engine will add the percentage values and the power values in the provided configuration fields. Error messages are issued for the following conditions:

- ▶ Power exceeds the **Max** value or falls below the **Min** value of the domain power.
- ▶ The **PercentageOfDomainBudget** values add up to over 100 percent.

Template:

```
{
  "@odata.context": "/redfish/v1/$Metadata#NvidiaNmDomain.NvidiaNmDomain",
  (continues on next page)
```

(continued from previous page)

```

"@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/0",
"@odata.type": "#NvidiaNmDomain.v1_4_0.NvidiaNmDomain",
"Capabilities": {
    "Max": 6000.0000,
    "Min": 4000.0000
},
"Id": "0",
"Name": "custom4",
"Status": {
    "State": "Enabled"
},
"Policies": {
    "@odata.context": "/redfish/v1/$Metadata#NvidiaNmPolicyCollection.
    ↪NvidiaNmPolicyCollection",
    "@odata.type": "#NvidiaNmPolicyCollection.NvidiaNmPolicyCollection",
    "Members": [
        {
            "@odata.context": "/redfish/v1/$Metadata#NvidiaNmPolicy.
            ↪NvidiaNmPolicy",
            "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/0/
            ↪Policies/0",
            "@odata.type": "#NvidiaNmPolicy.v1_2_0.NvidiaNmPolicy",
            "AssociatedDomainID": {
                "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/0"
            },
            "ComponentId": "COMP_CPU",
            "Id": "0",
            "Limit": 500.0000,
            "PercentageOfDomainBudget": 15.0000,
            "Name": "0"
        },
        {
            "@odata.context": "/redfish/v1/$Metadata#NvidiaNmPolicy.
            ↪NvidiaNmPolicy",
            "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/0/
            ↪Policies/1",
            "@odata.type": "#NvidiaNmPolicy.v1_2_0.NvidiaNmPolicy",
            "ComponentId": "COMP_MEMORY",
            "Id": "0",
            "Limit": 500.0000,
            "PercentageOfDomainBudget": 15.0000,
            "Name": "0"
        },
        {
            "@odata.context": "/redfish/v1/$Metadata#NvidiaNmPolicy.
            ↪NvidiaNmPolicy",
            "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/0/
            ↪Policies/2",
            "@odata.type": "#NvidiaNmPolicy.v1_2_0.NvidiaNmPolicy",
            "AssociatedDomainID": {
                "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/0"
            },
        }
    ]
}

```

(continues on next page)

(continued from previous page)

```

    "ComponentId": "COMP_GPU",
    "Id": "0",
    "Limit": 5000.0000,
    "PercentageOfDomainBudget": 70.0000,
    "Name": "0"
  }
],
"Members@odata.count": 3,
"Name": "NvidiaNmPolicyCollection"
}
}

```

- To create a new domain policy:

```
curl -k -u <bmc-user>:<password> -X POST https://<BMC>/redfish/v1/
  ↪Managers/BMC/NodeManager/Domains --data @<pathtojsonfile>
```

Example response:

```
{
  "@odata.context": "/redfish/v1/$Metadata#NvidiaNmDomain.NvidiaNmDomain",
  ↪",
  "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/Domains/14",
  "@odata.type": "#NvidiaNmDomain.v1_4_0.NvidiaNmDomain",
  "Capabilities": {
    "Max": 9000,
    "MaxCorrectionTimeInMs": 0,
    "MaxStatisticsReportingPeriod": "0",
    "Min": 6000,
    "MinCorrectionTimeInMs": 0,
    "MinStatisticsReportingPeriod": "0"
  },
  "Id": "14",
  "Name": "custom4",
  "Policies": {
    "@odata.context": "/redfish/v1/$Metadata#NvidiaNmPolicyCollection.
  ↪NvidiaNmPolicyCollection",
    "@odata.type": "#NvidiaNmPolicyCollection.NvidiaNmPolicyCollection
  ↪",
    "Members": [
      {
        "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/
  ↪Domains/14/Policies/0"
      },
      {
        "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/
  ↪Domains/14/Policies/1"
      },
      {
        "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/
  ↪Domains/14/Policies/2"
      }
    ]
  }
}
```

(continues on next page)

(continued from previous page)

```

        ],
        "Name": "NvidiaNmPolicyCollection"
    },
    "PolicyState": "Disabled",
    "Status": {
        "State": "Enabled"
    }
}

```

- ▶ To patch custom domain policies, provide only the configuration changes you want.
- ▶ Status codes for error responses:
 - ▶ 500: Indicates an error in the formatting of a particular field was incorrectly defined.
 - ▶ 400: Indicates that one or more fields were incorrect. In general, this will be for limits being out of bounds.

9.3.13.4 PSU Policies

Power supply unit (PSU) policies are read-only. PSU Policies set the overall available power budget for the system based on the number of active power supplies. The PSU Policy in effect enforces how the Domain Policies are selected.

- ▶ To view a list of PSU policies:

```
curl -k -u <bmc-user>:<password> https://<bmcip>/redfish/v1/Managers/BMC/
  ↳NodeManager/PSUPolicies
```

Example response:

```
{
    "@odata.context": "/redfish/v1/$Metadata#NvidiaNmPSUPolicyCollection.
  ↳NvidiaNmPSUPolicyCollection",
    "@odata.id": "/redfish/v1/Managers/BMC/NvidiaNmPSUPolicyCollection",
    "@odata.type": "#NvidiaNmPSUPolicyCollection.
  ↳NvidiaNmPSUPolicyCollection",
    "Members": [
        {
            "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/
  ↳PSUPolicies/0"
        },
        {
            "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/
  ↳PSUPolicies/1"
        },
        {
            "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/
  ↳PSUPolicies/2"
        },
        {
            "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/
  ↳PSUPolicies/3"
        },
    ]
}
```

(continues on next page)

(continued from previous page)

```
{
  "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/
  ↪PSUPolicies/4"
  }
],
"Members@odata.count": 5,
"Name": "NvidiaNmPSUPolicyCollection"
}
```

- ▶ To view a PSU policy:

```
curl -k -u <bmc-user>:<password> https://<bmcip>/Managers/BMC/NodeManager/
  ↪PSUPolicies/<PSUPolicyID>
```

For example, to view PSU policy 0:

```
curl -k -u <bmc-user>:<password> https://<bmcip>/Managers/BMC/NodeManager/
  ↪PSUPolicies/0
```

Example response:

```
{
  "@odata.context": "/redfish/v1/$Metadata#NvidiaNmPSUPolicy.
  ↪NvidiaNmPSUPolicy",
  "@odata.id": "/redfish/v1/Managers/BMC/NodeManager/PSUPolicies/0",
  "@odata.type": "#NvidiaNmPSUPolicy.v1_2_0.NvidiaNmPSUPolicy",
  "Id": "0",
  "LimitMax": 6000,
  "MaxPSU": 2,
  "MinPSU": 2,
  "Name": "Limp",
  "Status": {
    "State": "Disabled"
  }
}
```

PSU policy 0 defines the number of PSUs and the power that will be allocated to the system with a maximum of two PSUs.

- ▶ To view a metrics report:

A metrics report captures all critical values related to the power behavior of the system.

Example request:

```
curl -k -u <bmc-user>:<password> https://<bmcip>/redfish/v1/
  ↪TelemetryService/MetricReports/NvidiaNMMetrics_0
```

Example response:

```
{
  "@odata.id": "/redfish/v1/TelemetryService/MetricReports/
  ↪NvidiaNMMetrics_0",
  "@odata.type": "#MetricReport.v1_4_2.MetricReport",
  "Id": "NvidiaNMMetrics_0",
```

(continues on next page)

(continued from previous page)

```

    "MetricReportDefinition": {
        "@odata.id": "/redfish/v1/TelemetryService/
        MetricReportDefinitions/NvidiaNMMetrics_0",
        "MetricProperties": []
    },
    "MetricValues": [
        {
            "MetricId": "Domain_Policy_Active",
            "MetricValue": "12",
            "Timestamp": "2024-10-25T16:05:23+00:00"
        },
        {
            "MetricId": "PSU_Active_Policy",
            "MetricValue": "4",
            "Timestamp": "2024-10-25T16:05:23+00:00"
        },
        {
            "MetricId": "PSU_Redundancy_Policy",
            "MetricValue": "0",
            "Timestamp": "2024-10-25T16:05:23+00:00"
        },
        {
            "MetricId": "dcPlatformPower_avg",
            "MetricValue": "2549.00",
            "Timestamp": "2024-10-25T16:05:23+00:00"
        },
        ...
        {
            "MetricId": "gpuPowerCapabilitiesMax_7",
            "MetricValue": "1000.00",
            "Timestamp": "2024-10-25T16:05:23+00:00"
        }
    ],
    "Name": "NvidiaNMMetrics_0"
}

```

Table 1: Definitions of Metrics

MetricId	Definition	Example Metric Value
PSU_Active_Policy	The current active PSU pol- icy correponding to /red- fish/v1/Managers/BMC/NodeManager/PSUPolic	0-N /red- fish/v1/Managers/BMC/NodeManager/PSUPolic policy>
Domain_Policy_Active	The current active domain pol- icy corresponding to /red- fish/v1/Managers/BMC/NodeManager/Domain	0-N /red- fish/v1/Managers/BMC/NodeManager/Domain Active>

continues on next page

Table 1 – continued from previous page

MetricId	Definition	Example Metric Value
dcPlatformPower_avg	Total DC Power for the Platform	2181.00
dcPlatformPowerDGX_avg	Total DC Power for the non gpu base board components	1444.00
dcPlatformPowerHGX_avg	Total DC Power for the GPU Base Board	736.00
dcPlatformEnergy	Total Platform Energy (need to review)	2181.00
dcPlatformPowerLimit1		0.00
dcPlatformPowerLimit2		0.00
PSU_Redundancy_Policy	Current Policy Active PSU Policy	0
FixPwrDGXAvg	Power for fixed components on non gpu base board (e.g. FANS, NVMe, etc). Excludes CPU and Memory	1005.00
FixPwrHGXAvg	Power for fixed components on GPU Base Board. Excludes GPU	222.00
FixPwrAverage	Total Fixed Value for Platform	1228.00
Avb1NoCPU	Number of CPU	2
Avb1NoGPU	Number of GPU	8
PSU_WORKING_CNT	Total Number of PSU	6
DIMM_Count_Total	Total Number of DIMMS	32
GPU_PWR_BRAKE	State of Power Break	0
GPU_PWR_PRSNT	Indicates GPU Based Board is powered on	1
CPU_PWR_UNIT	Intel PWR Unit for CPU Power	3
CPU_TIM_UNIT	Intel Time Unit for CPU Energy	10
CPU_ENERGY_UNIT	Intel Energy Unit for CPU	14
cpuPackagePower_avg_0	Average Power for CPU0	193
cpuEnergy_0	Energy for CPU 0	196.00
coreEfficiency_0	Core Efficiency for CPU 0	61671.00
cpuPackagePowerCapabilitiesMin_0	Power Capabilities MIN CPU 0	209
cpuPackagePowerCapabilitiesMax_0	Power Capabilities MAX CPU 0	350
cpuPackagePowerLimit1_0	CPU Power Limit 1	400.00
cpuPackagePowerLimit2_0	CPU Power Limit 2	400.00

continues on next page

Table 1 – continued from previous page

MetricId	Definition	Example Metric Value
prochotRatioCapabilitiesMin_0	PROC Hot Ratio Min Capabilities CPU 0 (Min Frequency)	500
prochotRatioCapabilitiesMax_0	PROC Hot Ratio Max Capabilities CPU 0 (Max Frequency allowed when PROC Hot Asserted)	2000
turboRatioCapabilitiesMin_0	Turbo Ratio Min Capabilities CPU 0 (Min Frequency)	500
turboRatioCapabilitiesMax_0	Turbo Ratio Max Capabilities CPU 0 (Max Frequency)	3800
CPU_PWR_UNIT	Intel PWR Unit for CPU Power	3
CPU_TIM_UNIT	Intel Time Unit for CPU Energy	10
CPU_ENERGY_UNIT	Intel Energy Unit for CPU	14
cpuPackagePower_avg_1	Average Power for CPU1	182
cpuEnergy_1	Energy for CPU 1	185.00
coreEfficiency_1	Core Efficiency for CPU 1	62203.00
cpuPackagePowerCapabilitiesMin_1	Power Capabilities MIN CPU 1	209
cpuPackagePowerCapabilitiesMax_1	Power Capabilities MAX CPU 1	350
cpuPackagePowerLimit1_1	CPU Power Limit 1	400.00
cpuPackagePowerLimit2_1	CPU Power Limit 2	400.00
prochotRatioCapabilitiesMin_1	PROC Hot Ratio Min Capabilities CPU 1 (Min Frequency)	500
prochotRatioCapabilitiesMax_1	PROC Hot Ratio Max Capabilities CPU 1 (Max Frequency allowed when PROC Hot Asserted)	2000
turboRatioCapabilitiesMin_1	Turbo Ratio Min Capabilities CPU 1 (Min Frequency)	500
turboRatioCapabilitiesMax_1	Turbo Ratio Max Capabilities CPU 1 (Max Frequency)	3800
DIMM_Count_Socket_0	Number of DIMMS Socket 0	16.00
dramPackagePowerCapabilitiesMax_0	DRAM Power Capabilities MIN Socket 0	35.00
dramPackagePowerCapabilitiesMin_0	DRAM Power Capabilities MAX Socket 0	0.00
dramEnergy_0	DRAM Energy Socket 0	30.00
dramPowerLimit_0	DRAM Power Limit Socket 0	300.00

continues on next page

Table 1 – continued from previous page

MetricId	Definition	Example Metric Value
dramPower_avg_0	DRAM Average Power Socket 0	30.00
DIMM_Count_Socket_1	Number of DIMMS Socket 1	16.00
dramPackagePowerCapa-bilitiesMax_1	DRAM Power Capabilities MIN Socket 1	35.00
dramPackagePowerCapa-bilitiesMin_1	DRAM Power Capabilities MAX Socket 1	0.00
dramEnergy_1	DRAM Energy Socket 1	34.00
dramPowerLimit_1	DRAM Power Limit Socket 1	300.00
dramPower_avg_1	DRAM Average Power Socket 1	36.00
gpuPower_avg_0	GPU 0 Average Power	63.00
gpuPowerLimit_0	GPU 0 Power Limit	700.00
gpuPowerCapabili-tiesMin_0	GPU 0 Min Power Limit	200.00
gpuPowerCapabilities-Max_0	GPU 0 Max Power Limit	700.00
gpuPower_avg_1	GPU 1 Average Power	65.00
gpuPowerLimit_1	GPU 1 Power Limit	700.00
gpuPowerCapabili-tiesMin_1	GPU 1 Min Power Limit	200.00
gpuPowerCapabilities-Max_1	GPU 1 Max Power Limit	700.00
gpuPower_avg_2	GPU 2 Average Power	65.00
gpuPowerLimit_2	GPU 2 Power Limit	700.00
gpuPowerCapabili-tiesMin_2	GPU 2 Min Power Limit	200.00
gpuPowerCapabilities-Max_2	GPU 2 Max Power Limit	700.00
gpuPower_avg_3	GPU 3 Average Power	63.00
gpuPowerLimit_3	GPU 3 Power Limit	700.00
gpuPowerCapabili-tiesMin_3	GPU 3 Min Power Limit	200.00
gpuPowerCapabilities-Max_3	GPU 3 Max Power Limit	700.00
gpuPower_avg_4	GPU 4 Average Power	63.00
gpuPowerLimit_4	GPU 4 Power Limit	700.00

continues on next page

Table 1 – continued from previous page

MetricId	Definition	Example Metric Value
gpuPowerCapabilitiesMin_4	GPU 4 Min Power Limit	200.00
gpuPowerCapabilitiesMax_4	GPU 4 Max Power Limit	700.00
gpuPower_avg_5	GPU 5 Average Power	64.00
gpuPowerLimit_5	GPU 5 Power Limit	700.00
gpuPowerCapabilitiesMin_5	GPU 5 Min Power Limit	200.00
gpuPowerCapabilitiesMax_5	GPU 5 Max Power Limit	700.00
gpuPower_avg_6	GPU 6 Average Power	66.00
gpuPowerLimit_6	GPU 6 Power Limit	700.00
gpuPowerCapabilitiesMin_6	GPU 6 Min Power Limit	200.00
gpuPowerCapabilitiesMax_6	GPU 6 Max Power Limit	700.00
gpuPower_avg_7	GPU 7 Average Power	64.00
gpuPowerLimit_7	GPU 7 Power Limit	700.00
gpuPowerCapabilitiesMin_7	GPU 7 Min Power Limit	200.00
gpuPowerCapabilitiesMax_7	GPU 7 Max Power Limit	700.00

Chapter 10. Safety

This section provides information about how to safely use the NVIDIA DGX™ B200 system.

10.1. Safety Information

To reduce the risk of bodily injury, electrical shock, fire, and equipment damage, read this document and observe all warnings and precautions in this guide before installing or maintaining your server product.

In the event of a conflict between the information in this document and information provided with the product or on the website for a particular product, the product documentation takes precedence.

Your server should be integrated and serviced only by technically qualified persons.

You must adhere to the guidelines in this guide and the assembly instructions in your server manuals to ensure and maintain compliance with existing product certifications and approvals. Use only the described, regulated components specified in this guide. Use of other products or components will void the UL Listing and other regulatory approvals of the product and may result in noncompliance with product regulations in the region(s) in which the product is sold.

10.2. Safety Warnings and Cautions

To avoid personal injury or property damage, before you begin installing the product, read, observe, and adhere to all the following safety instructions and information.

The following safety symbols may be used throughout the documentation and may be marked on the product and the product packaging.

Symbol	Description
CAUTION	Indicates the presence of a hazard that may cause minor personal injury or property damage if the CAUTION is ignored.
WARNING	Indicates the presence of a hazard that may result in serious personal injury if the WARNING is ignored.

Symbol	Description
	Indicates potential hazard if indicated information is ignored.
	Indicates shock hazards that result in serious injury or death if safety instructions are not followed.
	Indicates hot components or surfaces.
	Indicates do not touch fan blades, may result in injury.
	<ul style="list-style-type: none"> ▶ Shock hazard: The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords. ▶ High leakage current ground (earth) connection to the Power Supply is essential before connecting the supply.
	Recycle the battery.
	The rail racks are designed to carry only the weight of the server system. Do not use rail-mounted equipment as a workspace. Do not place additional load onto any rail-mounted equipment.

安全警告和注意事項

為了避免人身傷害或財產損失，在開始安裝產品之前，閱讀，觀察，並遵守所有的以下安全提示和信息。以下安全符號會在整個文件中使用，並且在產品和/或產品包裝標記。

Symbol	Meaning
CAUTION	如果危險存在而忽略 CAUTION, 表示可能導致輕微的人身傷害和財產損失。
WARNING	如果危險存在而忽略 Warning, 表示可能會導致嚴重的人身傷害。
	表示如果信息被忽略會有潛在危險。
	表示如果不遵照安全說明會導致嚴重傷害或死亡的觸電危險
	表示零件或表面有熱度。
	表示請勿碰觸風扇葉片，碰觸可能會導致受傷
	<p>觸電危險 - 產品可能配備多根電源線。要除去所有危險電壓，請斷開所有電源線。</p> <p>在連接電源之前，請必須先連接高漏電電流地線到電源供應器。</p>
	電池回收。
	導軌架被設計成只能承受伺服器系統的重量。不要使用軌道安裝設備為工作區。不要把額外的負載加到任何軌道安裝設備上。

10.3. Intended Application Uses

This product was evaluated as Information Technology Equipment (ITE), which may be installed in offices, schools, computer rooms, and similar commercial type locations.

The suitability of this product for other product categories and environments (such as medical, industrial, residential, alarm systems, and test equipment), other than an ITE application, may require further evaluation.

10.4. Site Selection

Here is some information about how to select the correct site for the DGX B200 system. Choose a site that is:

- ▶ Clean, dry, and free of airborne particles (other than normal room dust).
- ▶ Well-ventilated and away from sources of heat including direct sunlight and radiators.
- ▶ Away from sources of vibration or physical shock.
- ▶ In regions that are susceptible to electrical storms, we recommend you plug your system into a surge suppressor and disconnect telecommunication lines to your modem during an electrical storm.
- ▶ Provided with a properly grounded wall outlet.
- ▶ Provided with sufficient space to access the power supply cord(s), because they serve as the product's main power disconnect.

10.5. Equipment Handling Practices

To reduce the risk of personal injury or equipment damage, do the following:

- ▶ Conform to local occupational health and safety requirements when moving and lifting equipment.
- ▶ Use mechanical assistance or other suitable assistance when moving and lifting equipment.

10.6. Electrical Precautions

10.6.1. Power and Electrical Warnings

Caution

The power button, indicated by the stand-by power marking, DOES NOT completely turn off the system AC power; standby power is active whenever the system is plugged in. To remove power from system, you must unplug the AC power cord from the wall outlet. Make sure all AC power cords are unplugged before you open the chassis, or add or remove any non hot-plug components.

Do not attempt to modify or use an AC power cord if it is not the exact type required. A separate AC cord is required for each system power supply.

Some power supplies in servers use Neutral Pole Fusing. To avoid risk of shock use caution when working with power supplies that use Neutral Pole Fusing.

The power supply in this product contains no user-serviceable parts. Do not open the power supply. Hazardous voltage, current and energy levels are present inside the power supply. Return to manufacturer for servicing.

When replacing a hot-plug power supply, unplug the power cord to the power supply being replaced before removing it from the server.

To avoid risk of electric shock, turn off the server and disconnect the power cords, telecommunications systems, networks, and modems attached to the server before opening it.

10.6.2. Power Cord Warnings

Caution

To avoid electrical shock or fire, check the power cord(s) that will be used with the product as follows:

- ▶ Do not attempt to modify or use the AC power cord(s) if they are not the exact type required to fit into the grounded electrical outlets.
- ▶ The power cord(s) must meet the following criteria:
 - ▶ The power cord must have an electrical rating that is greater than that of the electrical current rating marked on the product.
 - ▶ The power cord must have safety ground pin or contact that is suitable for the electrical outlet.
 - ▶ The power supply cord(s) is/ are the main disconnect device to AC power. The socket outlet(s) must be near the equipment and readily accessible for disconnection.
 - ▶ The power supply cord(s) must be plugged into socket-outlet(s) that is /are provided with a suitable earth ground.

10.7. System Access Warnings

To avoid personal injury or property damage, the following safety instructions apply whenever accessing the inside of the product:

- ▶ Turn off all peripheral devices connected to this product.
- ▶ Turn off the system by pressing the power button to off.
- ▶ Disconnect the AC power by unplugging all AC power cords from the system or wall outlet.
- ▶ Disconnect all cables and telecommunication lines that are connected to the system.
- ▶ Retain all screws or other fasteners when removing access cover(s). Upon completion of accessing inside the product, refasten access cover with original screws or fasteners.
- ▶ Do not access the inside of the power supply. There are no serviceable parts in the power supply.
- ▶ Return to manufacturer for servicing.
- ▶ Power down the server and disconnect all power cords before adding or replacing any non hot-plug component.
- ▶ When replacing a hot-plug power supply, unplug the power cord to the power supply being replaced before removing the power supply from the server.

Caution

If the server has been running, any installed processor(s) and heat sink(s) may be hot. Unless you are adding or removing a hot-plug component, allow the system to cool before opening the covers. To avoid the possibility of coming into contact with hot component(s) during a hot-plug installation, be careful when removing or installing the hot-plug component(s).

Caution

To avoid injury do not contact moving fan blades. Your system is supplied with a guard over the fan, do not operate the system without the fan guard in place.

10.8. Rack Mount Warnings

The following installation guidelines are required by UL to maintain safety compliance when installing your system into a rack.

The equipment rack must be anchored to an unmovable support to prevent it from tipping when a server or piece of equipment is extended from it. The equipment rack must be installed according to the rack manufacturer's instructions.

Install equipment in the rack from the bottom up with the heaviest equipment at the bottom of the rack.

Extend only one piece of equipment from the rack at a time.

You are responsible for installing a main power disconnect for the entire rack unit. This main disconnect must be readily accessible, and it must be labeled as controlling power to the entire unit, not just to the server(s).

To avoid risk of potential electric shock, a proper safety ground must be implemented for the rack and each piece of equipment installed in it.

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 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 air flow required for 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 (for example, the use of power strips).

10.9. Electrostatic Discharge

Caution

ESD can damage drives, boards, and other parts. We recommend that you perform all procedures at an ESD workstation. If one is not available, provide some ESD protection by wearing an antistatic

wrist strap attached to chassis ground (any unpainted metal surface) on your server when handling parts.

Always handle boards carefully. They can be extremely sensitive to ESD. Hold boards only by their edges. After removing a board from its protective wrapper or from the server, place the board component side up on a grounded, static free surface. Use a conductive foam pad if available but not the board wrapper. Do not slide board over any surface.

10.10. Other Hazards

10.10.1. CALIFORNIA DEPARTMENT OF TOXIC SUBSTANCES CONTROL

Perchlorate Material - Lithium Coin/Button Cell battery. Please dispose of properly.

Special handling may apply. See www.dtsc.ca.gov/hazardouswaste/perchlorate.

10.10.2. NICKEL



NVIDIA Bezel. The bezel's decorative metal foam contains some nickel. The metal foam is not intended for direct and prolonged skin contact. Please use the handles to remove, attach or carry the bezel. While nickel exposure is unlikely to be a problem, you should be aware of the possibility in case you are susceptible to nickel-related reactions.

10.10.3. Battery Replacement

Caution

There is the danger of explosion if the battery is incorrectly replaced. When replacing the battery, use only the battery recommended by the equipment manufacturer.

Dispose of batteries according to local ordinances and regulations. Do not attempt to recharge a battery.

Do not attempt to disassemble, puncture, or otherwise damage a battery.

更換電池警告：

警告

更換不正確之電池型式會有爆炸的風險。 請依製造商說明書處理用過之電池。

10.10.4. Cooling and Airflow

 **Caution**

Carefully route cables as directed to minimize airflow blockage and cooling problems. For proper cooling and airflow, operate the system only with the chassis covers installed.

Operating the system without the covers in place can damage system parts. To install the covers:

- ▶ Check first to make sure you have not left loose tools or parts inside the system.
- ▶ Check that cables, add-in cards, and other components are properly installed.
- ▶ Attach the covers to the chassis according to the product instructions.

The equipment is intended for installation only in a Server Room/Computer Room where both these conditions apply:

- ▶ Access can only be gained by SERVICE PERSONS or by USERS who have been instructed about the reasons for the restrictions applied to the location and about any precautions that shall be taken.
- ▶ Access is through the use of a TOOL or lock and key, or other means of security, and is controlled by the authority responsible for the location.

Chapter 11. Compliance

The NVIDIA DGX™ B200 Server is compliant with the regulations listed in this section.

11.1. United States

Federal Communications Commission (FCC)

FCC Marking (Class A)

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 any interference that may cause undesired operation of the device.

NOTE: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

California Department of Toxic Substances Control

Perchlorate Material - Lithium Coin/Button Cell battery. Please dispose it properly.
special handling may apply. See www.dtsc.ca.gov/hazardouswaste/perchlorate.

11.2. United States/Canada

TÜV Rheinland of North America is accredited as a Nationally Recognized Testing Laboratory (NRTL), by OSHA (The Occupational Safety and Health Administration) in the United States, and as a Product Certification Body by SCC (Standards Council of Canada) in Canada. Refer to <https://www.tuv.com/usa/en/ctuvus-certification.html>

cTUVus Mark



11.3. Canada

Innovation, Science and Economic Development Canada (ISED) CAN ICES(A)/NMB(A)

The Class A digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulation.

Cet appareil numerique de la class A respecte toutes les exigences du Reglement sur le materiel brouilleur du Canada.

11.4. CE

European Conformity; Conformité Européenne (CE)



This is a Class A product. In a domestic environment, this product may cause radio frequency interference in which case the user may be required to take adequate measures.

This device bears the CE mark in accordance with Directive 2014/53/EU. This device complies with the following Directives:

- ▶ EMC Directive A, I.T.E Equipment.
- ▶ Low Voltage Directive for electrical safety.
- ▶ RoHS Directive for hazardous substances.
- ▶ Energy-related Products Directive (ErP).

For the full text of EU declaration of conformity, refer to <http://www.nvidia.com/support>.

A copy of the Declaration of Conformity to the essential requirements may be obtained directly from NVIDIA GmbH (Bavaria Towers – Blue Tower, Einsteinstrasse 172, D-81677 Munich, Germany).

11.5. Australia and New Zealand

Australian Communications and Media Authority



This product meets the applicable EMC requirements for Class A, I.T.E equipment.

11.6. Brazil

INMETRO



11.7. Japan

Voluntary Control Council for Interference (VCCI)



この装置は、クラスA機器です。この装置を住宅環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

VCCI - A

This is a Class A product.

In a domestic environment, this product may cause radio interference, in which case the user may be required to take corrective actions. VCCI-A.

2008年、日本における製品含有表示方法、JISC0950が公示されました。製造事業者は、2006年7月1日以降に販売される電気・電子機器の特定化学物質の含有に付きまして情報提供を義務付けられました。製品の部材表示に付きましては、以下をご覧ください。

A Japanese regulatory requirement, defined by specification JIS C 0950, 2008, mandates that manufacturers provide Material Content Declarations for certain categories of electronic products offered for sale after July 1, 2006.

To view the JIS C 0950 material declaration for this product, visit www.nvidia.com.

Japan RoHS Material Content Declaration

日本工業規格 JIS C 0950:2008 により、2006年7月1日以降に販売される特定分野の電気および電子機器について、製造者による含有物質の表示が義務付けられます。						
機器名称: P4387B Server						
主な分類	特定化学物質記号					
	Pb	Hg	Cd	Cr(VI)	PBB	PBDE
PCB ボード	0	0	0	0	0	0
パッシブ電子部品	除外項目	0	0	0	0	0
アクティブ電子部品	除外項目	0	0	0	0	0
プロセッサー	0	0	0	0	0	0
メモリ	0	0	0	0	0	0
機械部品及びファン	除外項目	0	0	0	0	0
ケーブル/コネクター	除外項目	0	0	0	0	0
はんだ付け材料	0	0	0	0	0	0
ラベル、その他消耗品	0	0	0	0	0	0
注:						
1. 「0」は、特定化学物質の含有率が日本工業規格 JIS C 0950:2008 に記載されている含有率基準値より低いことを示します。						
2. 「除外項目」は、特定化学物質が含有マークの除外項目に該当するため、特定化学物質について、日本工業規格 JIS C 0950:2008 に基づく含有マークの表示が不要であることを示します。						
3. 「0.1wt%超」または「0.01wt%超」は、特定化学物質の含有率が日本工業規格 JIS C 0950:2008 に記載されている含有率基準値を超えていることを示します。						

A Japanese regulatory requirement, defined by specification JIS C 0950: 2008, mandates that manufacturers provide Material Content Declarations for certain categories of electronic products offered for sale after July 1, 2006.

Product Model Number: P4387B Server

Major Classification	Symbols of Specified Chemical Substance					
	Pb	Hg	Cd	Cr(VI)	PBB	PBDE
PCB	0	0	0	0	0	0
Passive components	Exempt	0	0	0	0	0
Active components	Exempt	0	0	0	0	0
Processor	0	0	0	0	0	0
Memory	0	0	0	0	0	0
Mechanical parts and Fan	Exempt	0	0	0	0	0
Cables/Connectors	Exempt	0	0	0	0	0
Soldering material	0	0	0	0	0	0
Flux, Solder Paste, label and other consumable materials	0	0	0	0	0	0

Notes:

1. "0" indicates that the level of the specified chemical substance is less than the threshold level specified in the standard, JIS C 0950: 2008.
2. "Exempt" indicates that the specified chemical substance is exempt from marking and it is not required to display the marking for that specified chemical substance per the standard, JIS C 0950: 2008.
3. "Exceeding 0.1wt%" or "Exceeding 0.01wt%" is entered in the table if the level of the specified chemical substance exceeds the threshold level specified in the standard, JIS C 0950: 2008.

11.8. South Korea

Korea Certification (KC)



R-R-WT1-P4387B

이 기기는 업무용 환경에서 사용할 목적으로 적합성평가를 받은 기기로서 가정용 환경에서 사용하는 경우 전파간섭의 우려가 있습니다

11.9. China

China Compulsory Certificate

No certification is needed for China. The NVIDIA DGX B200 is a server with rated current over than 6A.

China RoHS Material Content Declaration



部件名称 Parts	有害物质 Hazardous Substances					
	铅 (Pb)	汞 (Hg)	镉 (Cd)	六价铬 (Cr(VI))	多溴联苯 (PBB)	多溴二苯醚 (PBDE)
	○	○	○	○	○	○
PCB 板 PCB	○	○	○	○	○	○
被动电子零件 Passive components	X	○	○	○	○	○
主动电子零件 Active components	X	○	○	○	○	○
处理器 Processor	○	○	○	○	○	○
内存 Memory	○	○	○	○	○	○
结构件以及风扇 Mechanical parts and Fan	X	○	○	○	○	○
线材/连接器 Cable/Connectors	X	○	○	○	○	○
焊接金属 Soldering material	○	○	○	○	○	○
助焊剂, 锡膏, 标签及耗材 Flux, Solder Paste, label and other consumable materials	○	○	○	○	○	○

本表格依据 SJ/T 11364-2014 的规定编制

The table according to SJ/T 11364

○: 表示该有害物质在该部件所有均质材料中的含量均在 GB/T 26572 标准规定的限量要求以下。

○: Indicates that this hazardous substance contained in all of the homogeneous materials for this

part is below the limit requirement in GB/T 26572.

X: 表示该有害物质至少在该部件的某一均质材料中的含量超出 GB/T 26572 标准规定的限量要求。

X: Indicates that this hazardous substance contained in at least one of the homogeneous materials used for this part is above the limit requirement in GB/T 26572.

此表中所有名称中含 “X” 的部件均符合欧盟 RoHS 立法。

All parts named in this table with an “X” are in compliance with the European Union’s RoHS Legislation.

注：环保使用期限的参考标识取决于产品正常工作的温度和湿度等条件

Note: The referenced Environmental Protection Use Period Marking was determined according to normal operating use conditions of the product such as temperature and humidity.

11.10. Taiwan

Bureau of Standards, Metrology & Inspection (BSMI)



R33088
RoHS

警告使用者：
此為甲類資訊技術設備，於居住環境中使用時，可能會造成射頻擾動，在此種情況下，使用者會被要求採取某些適當的對策

報驗義務人：

香港商輝達香港控股有限公司台灣分公司 - 統一編號：80022300

臺北市內湖區基湖路8號。

Taiwan RoHS Material Content Declaration

Taiwan RoHS Material Content Declaration

限用物質含有情況標示聲明書 Declaration of the presence condition of the Restricted Substances Marking						
單元 Parts	限用物質及其化學符號 Restricted substances and its chemical symbols					
	鉛 (Pb)	汞 (Hg)	鎘 (Cd)	六價鉻 (Cr ⁶⁺)	多溴聯 苯 (PBB)	多溴二苯醚 (PBDE)
PCB 板 PCB	o	o	o	o	o	o
被動電子零件 Passive components	-	o	o	o	o	o
主動電子零件 Active components	-	o	o	o	o	o
處理器 Processor	o	o	o	o	o	o
內存 Memory	o	o	o	o	o	o
結構件以及風扇 Mechanical parts and Fan	-	o	o	o	o	o
線材/連接器 Cables/Connectors	-	o	o	o	o	o
焊接金屬 Soldering material	o	o	o	o	o	o
助焊劑，錫膏，標籤及耗材 Flux, Solder Paste, label and other consumable materials	o	o	o	o	o	o
備考 1: o: 系指該限用物質未超出百分比含量基準值 Note 1: o: indicates that the percentage content of the restricted substance does not exceed the percentage of reference value of presence.						
備考 2: -: 系指該項限用物質為排除項目。 Note 2: -: indicates that the restricted substance corresponds to the exemption.						
此表中所有名稱中含 “-” 的部件均符合歐盟 RoHS 立法。 All parts named in this table with an “-” are in compliance with the European Union's RoHS Legislation.						
注：環保使用期限的參考標識取決與產品正常工作的溫度和濕度等條件 Note: The referenced Environmental Protection Use Period Marking was determined according to normal operating use conditions of the product such as temperature and humidity.						

11.11. Russia/Kazakhstan/Belarus

Customs Union Technical Regulations (CU TR)



This device complies with the technical regulations of the Customs Union (CU TR)

ТЕХНИЧЕСКИЙ РЕГЛАМЕНТ ТАМОЖЕННОГО СОЮЗА О безопасности низковольтного оборудования (TP TC 004/2011)

ТЕХНИЧЕСКИЙ РЕГЛАМЕНТ ТАМОЖЕННОГО СОЮЗА Электромагнитная совместимость технических средств (TP TC 020/2011)

Технический регламент Евразийского экономического союза “Об ограничении применения опасных веществ в изделиях электротехники и радиоэлектроники” (TP ЕАЭС 037/2016)

Federal Agency of Communication (FAC)

This device complies with the rules set forth by the Federal Agency of Communications and the Ministry of Communications and Mass Media.

A Federal Security Service notification has been filed.

11.12. Israel

SII

ודא שלמות ותקינות כבל החשמל והתקע אין להזכיר או להוציא את התקע מרשות החשמל ביד"ם רטבות. אין לפתח את המכשיר, במקורה של בעיה כלשהו יש לפנות למועדת השירותים הקרה. יש להרחיק את המכשיר ממלחים. במקרה של ריח מוזר, רעשים שמקורם במכשיר, יש לנתקו מיידית מרשות החשמל ולפנות למועדת שירות המכשיר מיועד לשימוש בתוך המבנה, ולא לשימוש חיצוני ולא לשימוש בסביבה לחיה. אין לחתוך, לשבור, ועקב את הקבל החשמל. אין להניח חפצים על הקבל החשמל או להניח לו להתחמם יתר על המידה, שכן עלול לגרום למק, דלקה או התחשמלות. יש להקפיד לחזק את התקן הניתן במצב תקין מוקן לשימוש. אזהרה: אין להחליף את כבל הדינה בתחליפים לא מקוריים, חיבור לקוי עלול לגרום להתחשמלות המשמש. בשימוש על כבל מאריך יש לוודא תקינות מילוי הארקה שבכבל.

11.13. India

Bureau of India Standards (BIS)

Authenticity may be verified by visiting the Bureau of Indian Standards website at <http://www.bis.gov.in>.

India RoHS Compliance Statement

This product, as well as its related consumables and spares, complies with the reduction in hazardous substances provisions of the "India E-waste (Management and Handling) Rule 2016".

It does not contain lead, mercury, hexavalent chromium, polybrominated biphenyls or polybrominated diphenyl ethers in concentrations exceeding 0.1 weight % and 0.01 weight % for cadmium, except for where allowed pursuant to the exemptions set in Schedule 2 of the Rule.

11.14. South Africa

South African Bureau of Standards (SABS)

This device complies with the following SABS Standards:

SANS 2332: 2017/CISPR 32:2015 SANS 2335:2018/ CISPR 35:2016

National Regulator of Compulsory Specification (NRCS)

This device complies with following standard under VC 8055:

SANS IEC 60950-1

11.15. Great Britain (England, Wales, and Scotland)

UK Conformity Assessed

This device complies with the following Regulations:

- ▶ SI 2016/1091: Electromagnetic Compatibility (EMC)
- ▶ SI 2016/1101: The Low Voltage Electrical Equipment (Safety)
- ▶ SI 2012/3032: The Restriction of the Use of Certain Hazardous Substances in Electrical and Electronic Equipment (As Amended)

A copy of the Declaration of Conformity to the essential requirements may be obtained directly from NVIDIA Ltd. (100 Brook Drive, 3rd Floor Green Park, Reading RG2 6UJ, United Kingdom)

Chapter 12. Third-Party License Notices

This NVIDIA product contains third party software that is being made available to you under their respective open source software licenses. Some of those licenses also require specific legal information to be included in the product. This section provides such information.

12.1. Micron msecli

The [msecli utility](#) is provided under the following terms:

Micron Technology, Inc. Software License Agreement PLEASE READ THIS LICENSE AGREEMENT (“AGREEMENT”) FROM MICRON TECHNOLOGY, INC. (“MTI”) CAREFULLY: BY INSTALLING, COPYING OR OTHERWISE USING THIS SOFTWARE AND ANY RELATED PRINTED MATERIALS (“SOFTWARE”), YOU ARE ACCEPTING AND AGREEING TO THE TERMS OF THIS AGREEMENT. IF YOU DO NOT AGREE WITH THE TERMS OF THIS AGREEMENT, DO NOT INSTALL THE SOFTWARE. LICENSE: MTI hereby grants to you the following rights: You may use and make one (1) backup copy the Software subject to the terms of this Agreement. You must maintain all copyright notices on all copies of the Software. You agree not to modify, adapt, decompile, reverse engineer, disassemble, or otherwise translate the Software. MTI may make changes to the Software at any time without notice to you. In addition MTI is under no obligation whatsoever to update, maintain, or provide new versions or other support for the Software. OWNERSHIP OF MATERIALS: You acknowledge and agree that the Software is proprietary property of MTI (and/or its licensors) and is protected by United States copyright law and international treaty provisions. Except as expressly provided herein, MTI does not grant any express or implied right to you under any patents, copyrights, trademarks, or trade secret information. You further acknowledge and agree that all right, title, and interest in and to the Software, including associated proprietary rights, are and shall remain with MTI (and/or its licensors). This Agreement does not convey to you an interest in or to the Software, but only a limited right to use and copy the Software in accordance with the terms of this Agreement. The Software is licensed to you and not sold.

DISCLAIMER OF WARRANTY:

THE SOFTWARE IS PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND. MTI EXPRESSLY DISCLAIMS ALL WARRANTIES EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, NONINFRINGEMENT OF THIRD PARTY RIGHTS, AND ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. MTI DOES NOT WARRANT THAT THE SOFTWARE WILL MEET YOUR REQUIREMENTS, OR THAT THE OPERATION OF THE SOFTWARE WILL BE UNINTERRUPTED OR ERROR-FREE. FURTHERMORE, MTI DOES NOT MAKE ANY REPRESENTATIONS REGARDING THE USE OR THE RESULTS OF THE USE OF THE SOFTWARE IN TERMS OF ITS CORRECTNESS, ACCURACY, RELIABILITY, OR OTHERWISE. THE ENTIRE RISK ARISING OUT OF USE OR PERFORMANCE OF THE SOFTWARE REMAINS WITH YOU. IN NO EVENT SHALL MTI, ITS AFFILIATED COMPANIES OR THEIR SUPPLIERS BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, INCIDENTAL, OR SPECIAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION, OR LOSS OF INFORMATION) ARISING OUT OF YOUR USE OF OR INABILITY TO USE THE SOFTWARE, EVEN IF MTI

HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Because some jurisdictions prohibit the exclusion or limitation of liability for consequential or incidental damages, the above limitation may not apply to you.

TERMINATION OF THIS LICENSE: MTI may terminate this license at any time if you are in breach of any of the terms of this Agreement. Upon termination, you will immediately destroy all copies the Software.

GENERAL: This Agreement constitutes the entire agreement between MTI and you regarding the subject matter hereof and supersedes all previous oral or written communications between the parties. This Agreement shall be governed by the laws of the State of Idaho without regard to its conflict of laws rules.

CONTACT: If you have any questions about the terms of this Agreement, please contact MTI's legal department at (208) 368-4500. By proceeding with the installation of the Software, you agree to the terms of this Agreement. You must agree to the terms in order to install and use the Software.

12.2. Mellanox (OFED)

MLNX_OFED <<http://www.mellanox.com/>> is provided under the following terms:

Copyright (c) 2006 Mellanox Technologies. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Chapter 13. Notices

13.1. Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation (“NVIDIA”) makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer (“Terms of Sale”). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer’s own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or

services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

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

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF 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.

13.2. Trademarks

NVIDIA, the NVIDIA logo, DGX, DGX-1, DGX-2, DGX A100, DGX H100, DGX H200, DGX B200, DGX Station, and DGX Station A100 are trademarks and/or registered trademarks of NVIDIA Corporation in the United States and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

©2022-2025, NVIDIA Corporation