



# **NVIDIA DGX GB200 NVL72**

Release Notes for the 1.3.2 GA Multi-Node System Stack Firmware and Software Components for NVIDIA DGX GB200 System

# Document History

Version	Date	Description of Change
01	June 9, 2025	1.0 GA Release
01.1	June 30, 2025	1.0 GA Release - Switch ERoT Version typo fixed.
01.2	July 29, 2025	1.2.2 GA Release
01.3	Sept 30, 2025	1.3 GA Release
01.4	November 25, 2025	1.3.1 GA Release
01.5	January 9 <sup>th</sup> , 2026	1.3.2 GA Release

# Table of Contents

Overview .....	5
NVIDIA GB200 NVL72 Features .....	7
Multi-Node System Software Stack Package Contents .....	9
Compute Tray Components .....	9
Host Software Components .....	9
HMC .....	10
BMC .....	10
GB200 Switch Tray Components .....	11
BMC+FPGA+EROT .....	11
SBIOS+EROT .....	11
CPLD .....	12
NVOS .....	12
PowerShelf FW .....	13
NVIDIA/gdrcopy .....	13
System libraries, Tools, and Drivers .....	14
Improvements .....	17
Release 1.3.2 .....	17
Release 1.3.1 GA .....	18
Known Issues .....	20

# List of Tables

[Table 1. MFT/Tools Firmware Package Contents](#)

[Table 2. Public Release Links Associated with this Release](#)

[Table 3. CuDNN](#)

[Table 4. NCCL](#)

[Table 5. NVSHMEM](#)

[Table 6. DLFW Containers](#)

[Table 7. TRT-LLM](#)

[Table 8. NV Bandwidth Tool](#)

[Table 9. Tested Compatibility](#)

---

# Overview

This document contains detailed information about the SW 1.3.2 GA release for the NVIDIA® DGX GB200 72x1 rack configuration with support for trunk links in software/firmware packages.

This document includes dependencies and instructions that are specific to this release, the versions provided as part of the release, and a list of known issues.

**Legal Disclaimer:** The System SW 1.3.2 GA Release is only for the target product DGX GB200 NVL72 and cannot be used for any other product.



---

# NVIDIA GB200 NVL72 Features

This section provides a list of the DGX GB200 NVL features in this release.

This release includes the following features:



**Note:** No new features were introduced in the 1.3.2 release.

- NVLink Domain
  - 36x2 and 72x1 GPUs
- Fabric Manager
  - Support for NVLink Multi-Node Deployment.
  - Peer-to-Peer Traffic routing/configuration.
  - Single default partition with all available GPUs.
  - User Partition support.
- EGM
- Multicast support
- Software High availability
- Trunk link access link support
- IMEX Service
  - IMEX Service Peer-to-Peer Memory Import/Export support.
  - IMEX Dynamic Reconnect support.
  - IMEX Multi-User/Multi-Job Isolation support.
  - IMEX Service multicast (NVLink Sharp) Import/Export support.
- Platform information for the compute and switch trays.
- User Partition Support
- NVIDIA CUDA®
  - CUDA Driver and multi-node NVLink network and Arm®.
  - CUDA Toolkit 13.0 Update 2
- Multi-Node CUDA
  - Unicast traffic and multicast.
  - OpenMPI will be supported with NVLink Intranode and MNNVL internode.
  - Multi-Node NvBandwith.
- NCCL
  - NCCL: All2All, AllReduce, and Unicast.
  - NVSHMEM
- Math Libs
  - cuBlas (included in CUDA Toolkit 13.0 Update 2).
- NVSM
  - 25.3.9

- DCGM
  - 4.4.1 release.



---

# Multi-Node System Software Stack Package Contents

The tables in this section provide the contents of the system software stack package for the NVIDIA DGX GB200 NVL72 Rack. Refer to the [NVIDIA DGX GB200 NVL72 Systems Rack Firmware Update Guide](#) for more information about installation, configuration, and workarounds.

## Compute Tray Components



**Note:** The following are the host software versions tested with DGX GB200 Release 1.3.2.

## Host Software Components

Component	Version
DOCA_Host	3.2.0-125000
MFT Tools	4.34.0-145
GPU Driver	580.105.08
IMEX	580.105.08
WinOF-2	25.10.26947
CUDA	13.0.2
NMX-M	85.1.1100
CX7	28.47.1026
BF3_BFB	3.2

Component	Version
Host OS	DGX OS 7.3.1

## HMC

Bundle file: nvfw\_HGX-GBX00\_0023\_251209.1.0\_custom\_prod-signed.fwpkg

Component	Version
CPLD	0.22
GPU PS	97.00.B9.00.76
GPU QS	97.00.B9.00.75
EROT	01.04.0031.0000_n04
HMC	GB200Nv1-25.06-A
FPGA	1.60
SBIOS	02.04.14

Component	Filename
Recovery Firmware	nvfw_HGX-GBX00_0023_251209.1.0_custom_recovery_prod-signed.fwpkg
CoRIM	HGX-GBX00_0023_251209.1.0_custom_prod-signed.corim
CRT	nvfw_GB200-P4975_fwpkg_prod_cert.crt
PEM	nvfw_GB200-P4975_fwpkg_prod_pub.pem

## BMC

Bundle file: nvfw\_DGX-GBX00\_0023\_250901.1.0\_custom\_prod-signed.fwpkg

Component	Version
EROT	01.04.0031.0000_n04
BMC	GB200Nv1-25.06-A

Component	Filename
CoRIM	DGX-GBX00_0023_250901.1.0_custom_prod-signed.corim
Recovery Firmware	nvfw_DGX-GBX00_0101_250901.1.0_custom_recovery_prod-signed.fwpkg

## GB200 Switch Tray Components

### BMC+FPGA+EROT

Bundle file: [nvfw\\_GB200-P4978\\_0004\\_251130.1.0\\_prod-signed.fwpkg](#)

Component	Version
EROT	01.04.0026.0000
BMC	88.0002.1333
FPGA	0.1A

Component	Filename
CoRIM	GB200-P4978_0004_251130.1.0_prod-signed.corim
Recovery Firmware	nvfw_GB200-P4978_0004_251130.1.0_recovery_prod-signed.fwpkg

### SBIOS+EROT

Bundle file: [nvfw\\_GB200-P4978\\_0006\\_251130.1.0\\_prod-signed.fwpkg](#)

Component	Version
EROT	01.04.0026.0000

Component	Version
SBIOS	0ACTV_00.01.022

Component	Filename
CoRIM	GB200-P4978_0006_251130.1.0_prod-signed.corim
Recovery Firmware	nvfw_GB200-P4978_0006_251130.1.0_recovery_prod-signed.fwpkg

## CPLD

Bundle file: nvfw\_GB200-P4978\_0007\_251130.1.2\_prod-signed.fwpkg

Component	Version
CPLD1	CPLD000370_REV0600
CPLD2	CPLD000377_REV1700
CPLD3	CPLD000373_REV1200
CPLD4	CPLD000390_REV0400
FUI	OPN 000545

## NVOS

Component	Version
NVOS	25.02.2522
SM	2025.03.16
NMX-C	1.3.2_2025-12-01_16-49
GFM	570.195.06
NMX-T	1.3.2
QM3	35.2014.1874

Component	Filename
CoRIM	GB200-P4978_0005_251130.1.0_prod-signed.corim

## PowerShelf FW

Bundle file:

Component	Version
Powershell PSC	0104
Powershell PMC	3.2.4

## NVIDIA/gdrcopy

NVIDIA/gdrcopy is a fast GPU memory copy library-based on NVIDIA GPUDirect RDMA technology.

GDRCOPY version 2.5 is a publicly available release.

- The code is released on ([NVIDIA/gdrcopy](#)).
- The prebuilt packages release is available on [Index of /compute/redis/gdrcopy](#).
- Refer to [Magnum IO GDRCopy](#) for more information.

Component	Version
Stars	1006
Language	C++

---

# System libraries, Tools, and Drivers

**Table 1. MFT/Tools Firmware Package Contents**

The following table provides the versions of the necessary MFT/tools to download the firmware bundle.

Component	Version	Location
DCGM	4.4.1	<a href="https://developer.nvidia.com/dcgm#Downloads">https://developer.nvidia.com/dcgm#Downloads</a>
NVSM	25.03.09	Part of Installed BCM ISO

**Table 2. Public Release Links Associated with this Release**

The following table provides a list of the public links associated with this release.

Packages	Drivers Version	Tool Kit Version	Link
CUDA Drivers (580/TRD4)	580.105.08	13.0.2	<a href="https://developer.nvidia.com/datacenter-driver-580-105-08-download-archive">https://developer.nvidia.com/datacenter-driver-580-105-08-download-archive</a>
CUDA Toolkit 13.0.2 Download.		13.0.2	<a href="https://developer.nvidia.com/cuda-13-0-2-download-archive">https://developer.nvidia.com/cuda-13-0-2-download-archive</a>
NVLINK and CX7/BF3/DOCA	All Versions		<a href="https://docs.nvidia.com/networking/networking-nvlink/index.html">https://docs.nvidia.com/networking/networking-nvlink/index.html</a>

**Table 3. cuDNN**

The following table provides the download information for the cuDNN release.

Component	Version	Location
cuDNN	<b>9.15.0</b>	<a href="#">NVIDIA cuDNN release notes</a>

**Table 4. NCCL**

The following table provides information about where the NCCL release can be downloaded and links to the latest documentation.

Component	Location
NCCL Homepage	<a href="#">NVIDIA Collective Communications Library (NCCL) Download Page</a>
PIP Wheels	<a href="#">NVIDIA Collective Communication Library (NCCL) Runtime</a>
Github release	<a href="#">NVIDIA NCCL GitHub release</a>
Devzone binaries Current release <b>2.28.9</b>	<a href="#">NVIDIA NCCL download page</a> (current release) <a href="#">NVIDIA NCCL legacy download page</a> (legacy releases)
Release notes	<a href="#">NCCL release notes 2.28.9.</a> Note: GB200 support is available from NCCL 2.27.5

**Table 5. NVSHMEM**

The following table provides information about where the NVSHMEM release can be downloaded and links to the latest documentation.

Component	Location
NVSHMEM Homepage	<a href="#">NVIDIA NVSHMEM Homepage</a>
Release <b>3.4.5</b>	<a href="#">NVIDIA NVSHMEM Downloads</a>
Legacy releases	<a href="#">NVIDIA NVSHMEM Archive of legacy releases</a>
Documentation	<ul style="list-style-type: none"> <li>• Release Notes/Installation Guide: <a href="#">NVIDIA NVSHMEM Release Notes</a></li> <li>• API Documentation: <a href="#">NVIDIA OpenSHMEM Library (NVSHMEM) Documentation</a></li> <li>• Best Practice Guide: <a href="#">NVIDIA NVSHMEM Best Practice Guide</a></li> </ul>

**Table 6. DLFW Containers**

The following table provides a list of the [DL Frameworks containers](#).

Component	Versions	Location
PyTorch	Latest tag from NGC catalog.	<a href="#">NGC catalog for PyTorch</a>
Tensor RT	Latest tag from NGC catalog.	<a href="#">NGC catalog for TensorRT</a>
JAX	Latest tag from NGC catalog.	<a href="#">NGC catalog for JAX</a>

**Table 7. TRT-LLM**

Component	Version	Location
TRT-LLM	1.0	<a href="#">TensorRT-LLM</a>
vLLM	25.10	N/A
FlashInfer	0.2.6	N/A

**Table 8. NV Bandwidth Tool**

The following table provides information about the new version of the Multi-Node NV Bandwidth, for which the binary is included in the PID posting.

Component	Version	Location
Multi-Node NV Bandwidth	0.9	NVOnline: <b>1142140</b>

**Table 9. Tested Compatibility**

The following table lists the supported upgrade and downgrade paths for switch and compute components in this release.

Upgrade / Downgrade
<ul style="list-style-type: none"> <li>• 1.2.2 ↔ 1.3.0</li> <li>• 1.3.0 ↔ 1.3.2</li> <li>• 1.3.1 ↔ 1.3.2</li> </ul>



---

# Improvements

This section describes the improvements in each release.

## Release 1.3.2

1. **NVLink recovery is enabled by default in this release (NVLink improvement).**
2. **Duplicate NVLink plane ID leading to Isolated Xid 145 or NCCL all\_reduce\_perf failures with "uncorrectable NVLink error".**

**Workaround:** In earlier versions, use the following steps as a workaround:

1. Stop the NMX-C service: `nv action stop cluster app nmx-controller`
  2. Clear the Session Manager cache by removing files from `/var/log/nmx/nmx-c/nvlsn`: `rm -rf /var/log/nmx/nmx-c/nvlsn/*`
  3. Restart NMX-C service: `nv action start cluster app nmx-controller`
3. **Intermittent CPU/GPU throttling and degraded performance.**

This release includes filtering to prevent false thermal triggers, eliminating throttling and maintaining expected performance levels
  4. **LPDDR MultiBit ECC not in CMET.**

Fixed an issue where disabled memory channels could become re-enabled during CMET table updates.
  5. **Incorrect / out of range fan speed reported from NVSwitch.**

Resolved NVSwitch fan speed reporting issue caused by repeated CPLD inventory reads via BMC interface.
  6. **NVSwitch SSD was 100% used for syslog.**

Resolved an issue where syslog could fill the `/var/log` partition on GB200, causing 100% disk usage and system health alarms.
  7. **NVSwitch BMC \$expand >2 returns 404 on blacklisted Ports.**

Improved Redfish \$expand handling for levels >2 to gracefully skip or ignore responses from blacklisted/unsupported endpoints (for example, Switch Ports), eliminating spurious 404 Not Found errors while preserving complete data for supported resources.
  8. **Improved NVLink reliability with updated NVLink Reduction (NVR) and accounting for HLL (Head-of-queue Lifetime Limit) when NVLink Recovery is enabled.**
  9. **Switch detected as unhealthy but the partition was never marked as unhealthy.**

This fix identifies the stale GPU health state after the compute tray removal and marks the GPU as GPU\_HEALTH\_NO\_NVLINK.
  10. **Fixed intermittent compute node hangs during heavy multicast object operations.**

Fixed a regression that caused intermittent compute node hangs during heavy multicast object churn, ensuring multicast-intensive workloads complete reliably under full cluster load.

## Release 1.3.1 GA

### 1. **Intermittent C2C link training failure during GPU reset.**

During GPU reset, the Grace<->GPU C2C link undergoes training. In certain cases, this process intermittently fails, leading to unrecoverable host errors and, in some instances, Grace firmware crashes causing the host to reboot. This issue has been resolved in the updated GPU firmware.

### 2. **In rare instances, the PKEY request is rejected by Subnet Manager.**

Fixed a race condition where the PKEY database is always cleared before sending the response to the GFM.

### 3. **Resource leak in RDM upon multicas join cancellation.**

Multicast join request waiting in the RDM queue were cancelled by the GFM leading to a resource leak. Fix ensures that cancelled requests are properly removed from the RDM queue, preventing unnecessary resource allocation and eliminating the leak.

### 4. **FLR reset hangs during GPU engine transaction.**

A FLR reset, initiated by the GPU FLR handler could result in a hang. Added a fix to leverage the hardware engine reset state machine to reliably reset the engine.

### 5. **Failed to find an active link to send remap tables.**

Removing GPU from the topology, caused Fabric Manager to send configuration updates to GPUs over invalid links, resulting in errors. Added a fix to ensure that link down events are properly propagated, allowing Fabric Manager to use refreshed and accurate link information when sending configuration updates.

### 6. **Fixed multicast reference count computation.**

Fixed an issue that caused multicast reference count to be computed incorrectly.

### 7. **Allow GPUs that are part of a multicast group to be removed from the partition.**

Added a fix to allow GPUs that are part of a multicast group to be removed from a partition.

### 8. **Added a GFM config option to control the automatic reroute of unhealthy partitions.**

Automatic reroute of unhealthy partitions at NMX-C startup is controlled using a config option `MNNVL_AUTO_REROUTE_UNHEALTHY_PARTITION_ON_INIT` which is set to true by default.

### 9. **Partitions are deleted from the domain.**

Resolved a crash issue in the GFM caused by logging variables with non-null terminated strings.

**10. GFM is responding to older versions of Probe Requests.**

In rare scenarios, the GFM might respond to older versions of probe requests from the GPU driver, causing incorrect remap programming and impacting traffic. Added a fix to ignore older versions of probe requests resulting in the GPU driver sending a follow-up probe request with the correct version.

**11. Fixed GPU health reporting when a GPU is reset or the compute tray is rebooted.**

When a GPU is reset or the compute tray is rebooted, the GPU health can incorrectly reported as HEALTHY. Added a fix to report the GPU health as NO\_NVLINK indicating that the GPU can't participate in NVLink traffic in this state.

**12. ETH0/1 routing was not separated, causing incorrect routing when eth0 was disconnected.**

Fixed with separate routing per interface (eth0/1).

**13. Improved auto-reroute of trunk links when partition is unhealthy after NMX-C or hosting switch restart.**

With this fix, if a partition is detected as unhealthy due to a trunk link routing issue, restarting NMX-C or rebooting the hosting switch now automatically attempts to reroute trunk links and restore the partition to a healthy state, provided no trunk link issues persist.

**14. Three different threads in GFM had a circular dependency on the locks, resulting in deadlock.**

Added a fix that removes conditions for the deadlock.

**15. NVLS bind failures occur.**

Implemented fixes across the GPU driver, Fabric Manager, and GPU firmware to improve the reliability of NVLS initialization.

---

# Known Issues

This section provides a list of the known issues.

1. **On certain DGX GB200 NVSwitch platforms, the switch reboot history may display multiple unexpected reboots with the reason shown as Platform reset (retrieved using the NVOS command: `nv show sys reboot -o json | jq -r .reason`). NVIDIA plans to include a fix in a future NVOS release.**

2. **NVIDIA compute firmware recovery bundle fails integrity check.**

In release 1.3.2, updating GPU firmware using the recovery bundle may fail with the Integrity check failed for FW package message in the Firmware Update Service. The update is rejected before the image is applied, and the GPU remains on its previous firmware version.

**Workaround:** Use the standard compute-tray firmware bundle for GB200 NVL72 v1.3.2 (non-recovery ...\_prod-signed.fwpkg) to update or restore GPU firmware. If a recovery bundle is required, contact NVIDIA to obtain a corrected package and retry the update with that bundle.

3. **NETIR dumps missing, libnvidia-ml.so error in NV-Bug-Report**

**Workaround:** Use the following steps as a temporary workaround. The symbolic link will be restored in a future release.

1. Use one of the following commands to locate the **libnvidia-ml.so.1** file.  

```
find /usr -name "libnvidia-ml.so.1"  
locate libnvidia-ml.so.1
```
2. Use one of the following commands to create a symbolic link named **libnvidia-ml.so** pointing to the located file in one of the standard library installation paths.  

```
ln -sf <path-to-libnvidia-ml.so.1> /usr/lib64/libnvidia-ml.so  
ln -sf <path-to-libnvidia-ml.so.1> /usr/lib/libnvidia-ml.so
```

4. **2+ domain NCCL All-to-All testing may result in system or deadlock hangs with no error codes.**

**Workaround:** There is currently no workaround, and this issue will be fixed in a future release. Future versions of NCCL include atomic locks to code path to prevent the race condition.

5. **A GPU that is part of a multicast group cannot be removed from the partition used in MultiCast can be reused prematurely.**

In GB200/GB300, if a GPU is participating in an active Multicast Team, it cannot be removed from a partition.

**Workaround:**

- Reset or power cycle all partition GPUs – On startup, the probe request will release any previously allocated multicast groups. Once the probe request completes, the GPU can then be removed from the partition.
- Delete and recreate the partition – Delete the current partition and create a new one that excludes the GPUs intended for removal. NVIDIA recommends using a different partition ID for the new partition.

**6. GPU reset operations may fail with kernel Xid errors (for example, GSP Timeout) and require a rack-level power cycle to recover if a subset of NVSwitch trays are rebooted, either unexpectedly or through an orderly reboot, while an NVLink Sharp workload is actively running.**

**Workaround:** If the GPU encounters this issue, a system power cycle is required to restore functionality. Either an AC or DC power cycle will bring the GPU back online.

**7. Occasionally, after a GPU reset, the Fabric status in nvidia-smi/NVML may report Insufficient Resources, even though all GPU NVLinks are active.**

**Workaround:** Reset the affected GPUs again. If no GPU-resident services (such as nvidia-persistenced) are running, allow a sufficient amount of time (~10 seconds) between the GPU reset and checking the Fabric status.

**8. NVLink NVLS traffic fails with Xid 145 after partial switch tray reboot**

**Workaround:** If only a subset of switch trays are rebooted, perform an additional NMX-C restart after all switch trays are back online. Or, avoid rebooting a subset of switches and instead, reboot/power cycle all switch trays.

**9. Micron NVMe drives may become unavailable when connected behind the Mellanox CX8 due to a PCIe Max Payload Issue. This issue does not impact GB200 with CX7, but it does impact GB200 with CX8.**

**Workaround:** Set the PCIe Max Payload setting to 128bytes in the SBIOS setup. A firmware patch in the SBIOS is also available but currently not included in the SBIOS in the 1.3.0 bundle.

**10. Type 3 Command 0x4A GPM Query per-instance GPM metrics: Querying some metrics will get 0x7e ERR\_BUSY.**

**Workaround:** There is currently no workaround, and this issue will be fixed in a future firmware release.

**11. Type 3 Command 0x02 Read Thermal Parameter: unknown Target GPU TLIMIT temperature.**

**Workaround:** There is currently no workaround. The target T.Limit threshold does not apply to this product, and the value can be safely ignored.

**12. Type 3 Command 0x05 Clear Max Observed Power: Got ERR\_INVALID\_DATA when sensor ID=255.**

**Workaround:** Specify the sensors to be cleared one at time instead of clearing them all at once using the aggregate sensor ID (255).

**13. The Grace RW SPI read functionality needs to be improved to meet the RMA criteria.**

Based on internal testing, the read full SPI time to read is one hour.

**Workaround:** There is currently no workaround, and this issue will be fixed in a future release.

**14. The NVSwitch EROT recovery is not working.**

**Workaround:** Customers can still recover a failed EROT and can use the typical normal firmware update mechanisms to update EROTs.

**15. BMC PCIe link reset causes SBIOS exception.**

**Workaround:** Customers should avoid ungraceful PCIe link resets to the BMC while systems are operational. Partners that do not route PCIe to the BMC are not impacted by this issue.

**16. coRIM should reference the value of FWID tcg-dice-TcbInfo in AliasKeyCert.**

The **FWID[0]** field of `DiceTcbInfo` in the GPU iRoT DICE certificate contains the FSP FMC measurement. This value is written to an RTS hardware register (MSR 2), with hardware recording additional state information to compute the value to be stored in the register. The final register value is reported in SPDM measurement block Index 4.

The following Python code shows how to generate the measurement block 4 value in the GPU CoMID based from the FWID value contained in the DICE certificate.

**17. Firmware Update task completed 100% but says critical.**

A full log file causes the task status update to fail.

**Workaround**

The fw update was successful, no workaround is needed. To see the status switch to successful, clear the eMMC and run it again:

```
https://${bmcip}/redfish/v1/Managers/HGX_BMC_0/Actions/Oem/eMMC.SecureErase
```

**18. GPU firmware cannot be updated until after SBIOS boots into UEFI.**

The Grace PCIe PERST signal causes the GPU firmware's MCTP stack to not work reliably until the UEFI stage of boot (PERST de-asserted).

**Workaround**

Wait for the SBIOS to boot into UEFI to run a firmware update.

**19. The SBIOS will not boot on after an L2 reset. ☐☐**

If the Boot Chain is corrupt or does not boot, the Grace/EROT tries to fall back to Boot Chain 1. On an earlier version of the SBIOS, Socket 0 successfully resets, but Socket 1 fails with a SPI contention.

**Workaround**

Ensure that Boot Chain 0 contains a proper image.

## **20. After an Aux or PDU power cycle, the BMC may fail to restore the host to previous power state**

If PowerRestorePolicy is set to LastState, the BMC may occasionally fail to restore to the previous host power state after an Aux or PDU power cycle. For example, if the previous host state was On, the BMC may not turn on the host automatically after power is restored.

### **Workaround**

If there is a fixed desired power state for the host after the power is restored, user can set the PowerRestorePolicy to either AlwaysOn or AlwaysOff. Otherwise, the user can send an additional power on or power off command after the power is restored to achieve the desired host power state.

## **21. Leak Detection Voltage Sensors does not show Degraded State**

When the BMC gets an out-of-range reading from the voltage-based leak detection sensor, it marks the Leak Detector's Status:State as Degraded. The corresponding Voltage Sensor does not currently align its Status:State property with the Leak Detector to reflect the Degraded value.

### **Workaround**

It is recommended to use the Leak Detector Redfish URIs to get the most accurate state of the leak detector system. The sensor will only provide the raw voltage value for diagnostic purposes.

Leak Detector URIs:

```
/redfish/v1/Chassis/Chassis_0/ThermalSubsystem/LeakDetection/LeakDetectors/Chassis_0_LeakDetector_0_ColdPlate
```

```
/redfish/v1/Chassis/Chassis_0/ThermalSubsystem/LeakDetection/LeakDetectors/Chassis_0_LeakDetector_0_Manifold
```

```
/redfish/v1/Chassis/Chassis_0/ThermalSubsystem/LeakDetection/LeakDetectors/Chassis_0_LeakDetector_1_ColdPlate
```

```
/redfish/v1/Chassis/Chassis_0/ThermalSubsystem/LeakDetection/LeakDetectors/Chassis_0_LeakDetector_1_Manifold
```

## **22. No mechanism to disable host IPMI interface**

This firmware release does not provide a mechanism to disable the host IPMI interface and any privileged user running on a compute tray can send IPMI commands to the BMC on the same compute tray.

### **Workaround**

A future BMC firmware release will provide an interface to restrict host IPMI commands. In the meantime, necessary access controls need to be implemented on the compute trays to limit host privileged access by users not authorized to access the BMC.

### **23. BMC SKU ID is sometimes not available through Redfish**

The BMC's SKU ID should be available through the following Redfish URI under the "SKU" property:

```
/redfish/v1/Chassis/BMC_0
```

Due to a current issue on the BMC, this property may not be available on all systems.

#### **Workaround**

There is currently no known workaround. The plan is to resolve this issue for the next release.

### **24. Can not set IPv6 address through Redfish**

There is a current BMC bug where setting the IPv6 through the following Redfish URI fails:

```
/redfish/v1/Managers/BMC_0/NetworkProtocol
```

The command will return the following message:

```
"Message": "The property 'Address or Port' with the requested value of  
'\"2620:10d:c0a3:103::9:7554\"' could not be written because the value does not meet the  
constraints of the implementation.",
```

This is a known issue and will be addressed in the next release.

#### **Workaround**

There is currently no known workaround. The plan is to resolve this issue for the next release.

### **25. No error Redfish response when credentials are not provided**

If no credentials are provided in a Redfish command, the BMC currently does not return an explicit message indicating the lack of credentials. A **401** Unauthorized response is still returned, and access is prevented. This only applies if no credentials are provided. If credentials are provided but are incorrect, then the BMC will return a message indicating invalid credentials.

#### **Workaround**

The return code can be checked to confirm the status of the command. The **401** will indicate to the user that access was unauthorized.

### **26. Redfish GPU\_0/Ports/NVLink\_0, the property "LinkState" value does not change after a DC Cycle.**

NVLinks that were previously disabled through Redfish continue to report the LinkState as Enabled and LinkStatus as LinkUp.

#### **Workaround**



Use host reporting tools (nvidia-smi/NVML/DCGM) to fetch the NVLink status for disabled links.

## **27. No option in the webgui to power on the system**

Currently, There is no option to power on the System from the BMC Webgui.

### **Workaround**

Use redfish or ipmitool to power on the system.

## Notice

The information provided in this specification is believed to be accurate and reliable as of the date provided. However, NVIDIA Corporation ("NVIDIA") does not give any representations or warranties, expressed or implied, as to the accuracy or completeness of such information. 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 publication supersedes and replaces all other specifications for the product that may have been previously supplied.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and other changes to this specification, at any time and/or to discontinue any product or service without notice. Customer should obtain the latest relevant specification 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. NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this specification.

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 these specifications will be suitable for any specified use without further testing or modification. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to ensure the product is suitable and fit for the application planned by customer and to do 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 specification. NVIDIA does not accept any 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 specification, 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 specification. 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 specification is permissible only if reproduction is approved by NVIDIA in writing, is reproduced without alteration, and is accompanied by all associated conditions, limitations, and notices.

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. 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 NVIDIA terms and conditions of sale for the product.

## Trademarks

NVIDIA, the NVIDIA logo, HGX, NVSwitch, and NVLink are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company product names may be trademarks of the respective companies with which they are associated.

## Copyright

© 2025 NVIDIA CORPORATION & AFFILIATES. All rights reserved.