# NVIDIA DGX GB300 NVL72

## Release Notes v1.0.1 for Multi-Node System Stack Firmware and Software

RN-11874-001_1.0.1 | November 2025

# Document History

RN-11874-001_1.0.1

| Version | Date | Description of Change |
|---|---|---|
| 01 | Oct 10, 2025 | Release 1.0.0 |
| 02 | Dec 19, 2025 | Release 1.0.1<br>• Added improvements and known issues |

# Table of Contents

# List of Tables

# Overview

This document contains detailed information for the NVIDIA® DGX GB300 NVL72 v1.0.1 software/firmware release.

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 and improvements.

**Legal Disclaimer:** The System SW 1.0.1 GA Release is only for the target product DGX GB300 NVL72 and cannot be used for any other product.

# NVIDIA DGX GB300 NVL72 Features

This section outlines the DGX GB300 NVL72 features included in this release, which provides essential operational capabilities to support end-to-end system qualification and validation.

This release includes the following features:

- NVLink Domain
  - 72x1 GPUs
- Fabric Manager
  - Support for NVLink multi-node deployment.
  - Peer-to-Peer Traffic routing/configuration.
  - Single default partition and multiple user partitions, with all available GPUs.
- IMEX Service
  - IMEX Service Peer-to-Peer Memory Import/Export support.
  - IMEX Dynamic Reconnect support.
  - IMEX Service multicast (NVLink SHARP) Import/Export support.
- RM Multi-User/Multi-Job Isolation support.
- Platform information for the compute and switch trays.
- NVIDIA multi-node CUDA®
  - Unicast traffic and multicast.
  - OpenMPI will be supported with NVLink Intranode and MNNVL internode.
  - Multi-Node NvBandwidth.
- NCCL
  - NCCL: All2All, AllReduce, and Unicast.
  - NVSHMEM
- nvidia.nvlink Ansible Collection
  - v1.6.9
- CX8
  - CX8 Telemetry
- DCGM with multi-node testing support

# 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 GB300 NVL72 reference design. Refer to the *FW update guide* for more information about installation, configuration, and workarounds.

This release includes public posting links necessary for creating the Multi-Node systems stack.

# Compute Tray Components

Refer to the *NVIDIA FW update guide* for more information about the HMC and HMC NO-BIOS firmware bundle and installation instructions.

**Note:** The following are the host software versions tested with GB300 Release 1.0.1.

## 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 |
| Kernel Module Source | NVIDIA-kernel-module-source-580.105.08.tar.xz |
| CUDA | 13.0.1 |
| WinOF-2 | 25.10.50000 |
| NMX-M | 85.1.1100 |
| BF3_BFB | 3.2.0-113 |

## HMC

Bundle file: nvfw_HGX-GBX00_0026_251126.1.0_custom_prod-signed.fwpkg

| Component | Version |
|---|---|
| CPLD | 0.22 |
| GPU | 97.10.4A.00.0C |

| Component | Version |
|---|---|
| EROT | 01.04.0031.0000_n04 |
| HMC | GB200Nvl-25.08-B |
| SBIOS | 02.05.13 |
| FPGA | 1.60 |

| Component | Filename |
|---|---|
| CoRIM | HGX-GBX00_0026_251126.1.0_custom_prod-signed.corim |
| Recovery Firmware | nvfw_HGX-GBX00_0026_251126.1.0_custom_recovery_prod-signed.fwpkg |
| CRT | nvfw_GB300-P4059_fwpkg_prod_cert.crt |
| PEM | nvfw_GB300-P4059_fwpkg_prod_pub.pem |
| VBIOS Only Firmware | g548_0301_893_97104A000C.fwpkg |

# BMC

Bundle file: nvfw_DGX-GBX00_0026_251112.1.0_custom_prod-signed.fwpkg

| Component | Version |
|---|---|
| EROT | 01.04.0031.0000_n04 |
| BMC | GB200Nvl-25.08-9 |
| SMA Firmware | 0003.00.0220.0001 |

| Component | Filename |
|---|---|
| CoRIM | DGX-GBX00_0026_251112.1.0_custom_prod-signed.corim |
| CRT | nvfw_GB300-p4058_corim_prod_cert.crt |
| PEM | nvfw_GB300-p4058_corim_prod_pub.pem |

# CX8

| Component | Filename |
|---|---|
| SP0 Firmware | fw-ConnectX8-rel-40_47_1026-900-9X86E-00CX-SP0_Ax-UEFI-14.40.10-FlexBoot-3.8.201.signed-MT_0000001226.fwpkg |
| SP0 CoRIM | fw-ConnectX8-rel-40_47_1026-900-9X86E-00CX-SP0_Ax-UEFI-14.40.10-FlexBoot-3.8.201.signed.signed.corim |
| ST0 Firmware | fw-ConnectX8-rel-40_47_1026-900-9X86E-00CX-ST0_Ax-UEFI-14.40.10-FlexBoot-3.8.201.signed-MT_0000001513.fwpkg |
| ST0 CoRIM | fw-ConnectX8-rel-40_47_1026-900-9X86E-00CX-ST0_Ax-UEFI-14.40.10-FlexBoot-3.8.201.signed.signed.corim |

# GB300 PS Switch Tray Components

## NVOS

| Component | Version |
|---|---|
| NVOS | 25.02.4317 |
| SM | 2025.06.12 |
| GFM | R580.109 |
| NMX-C | 4.21.30 |
| NMX-T | 3.5.1 |
| QM3 | 35.2014.4712 |

| Component | Filename |
|---|---|
| CoRIM | GB300-P4093_0005_251112.1.1_prod-signed.corim |

## BMC+FPGA+EROT

Bundle file: nvfw_GB300-P4093_0004_251112.1.1_prod-signed.fwpkg

| Component | Version |
|---|---|
| EROT | 01.04.0031.0000_n04 |
| BMC | 88.0002.1958 |
| FPGA | 0.24 |

| Component | Filename |
|---|---|
| CoRIM | GB300-P4093_0004_251112.1.1_prod-signed.corim |

| Component | Filename |
|---|---|
| Recovery Firmware | nvfw_GB300-P4093_0004_251112.1.1_recovery_prod-signed.fwpkg |

# SBIOS+EROT

Bundle file: nvfw_GB300-P4093_0006_251112.1.1_prod-signed.fwpkg

| Component | Version |
|---|---|
| EROT | 01.04.0031.0000_n04 |
| SBIOS | 0ACTV_00.01.020 |

| Component | Filename |
|---|---|
| CoRIM | GB300-P4093_0006_251112.1.1_prod-signed.corim |
| Recovery Firmware | nvfw_GB300-P4093_0006_251112.1.1_recovery_prod-signed.fwpkg |

# CPLD

Bundle file: nvfw_GB300-P4093_0007_251112.1.1_prod-signed.fwpkg

| Component | Version |
|---|---|
| CPLD1 | CPLD000420_REV0300 |
| CPLD2 | CPLD000419_REV0301 |
| CPLD3 | CPLD000418_REV0200 |

# 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/redist/gdrcopy.
- Refer to Magnum IO GDRCopy for more information.

| Component | Version |
|---|---|
| Stars | 1006 |
| Language | C++ |

# System Software, Tools, and Drivers

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

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

| Description | Link |
|---|---|
| Datacenter Driver Version **580.105.08** | Datacenter Driver 580 downloads |
| CUDA Toolkit **13.0.2** | CUDA Toolkit release<br>**Note**: This is the link for the current release which might not match 13.0.2. If the current release does not match, locate release13.0.2 in the archives at Archive of previous CUDA releases. |

### Table 3. cuDNN

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

| Component | Version | Location |
|---|---|---|
| cuDNN | **9.16.0** | NVIDIA cuDNN release notes |

### 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 | NVIDIA Collective Communications Library (NCCL) Download Page |
| PIP Wheels | NVIDIA Collective Communication Library (NCCL) Runtime |
| Github release | NVIDIA NCCL GitHub release |
| Devzone binaries<br>Current release **2.28.7** | NVIDIA NCCL download page (current release)<br>NVIDIA NCCL legacy download page (Legacy releases) |
| Release notes **2.28.7** | NCCL release notes |

### Table 5. NVSHMEM

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

| Component | Location |
|---|---|
| NVSHMEM Homepage | NVIDIA NVSHMEM Homepage |
| Release **3.4.5** | NVIDIA NVSHMEM Downloads |
| Legacy releases | NVIDIA NVSHMEM Archive of legacy releases |
| Documentation | • Release Notes/Installation Guide: NVIDIA NVSHMEM Release Notes<br>• API Documentation: NVIDIA OpenSHMEM Library (NVSHMEM) Documentation<br>• Best Practice Guide: NVIDIA NVSHMEM Best Practice Guide |

### Table 6. DLFW Containers

The following table provides a list of the DL Frameworks containers.

| Component | Versions | Location |
|---|---|---|
| PyTorch | Pytorch version **2.8.0a0** with Nvidia additions as part of DLFW 25.08 container | NGC catalog for PyTorch |
| Tensor RT | TensorRT Version **10.13.2.2** | NGC catalog for TensorRT |
| JAX | JAX version **0.6.2** with Nvidia additions as a part of DLFW **25.08** Container | NGC catalog for JAX |
| DL Frameworks Release Notes | **25.08** | NVIDIA Optimized Frameworks - NVIDIA Docs |

### Table 7. NVIDIA TensorRT

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

**Attention**: TensorRT version 10.13 includes functional support for GB300. TensorRT (to be released at a future date) will offer performance improvements that utilize GB300 system capabilities.

| Component | Location |
|---|---|
| TensorRT Homepage | NVIDIA TensorRT Getting Started Documentation<br>Nvonline: **1139359, 1139351,1139347,1139356** |
| Release **10.14.1** | NVIDIA TensorRT current releases |
| Legacy releases | NVIDIA TensorRT legacy releases |
| Documentation | • Release Notes/Installation Guide: NVIDIA TensorRT Overview and Release Notes<br>• API Documentation: NVIDIA TensorRT Documentation |
| TRT-LLM | nvcr.io/nvidia/tensorrt-llm/release:gb300-ea<br>nvcr.io/nvidia/tensorrt-llm/devel:gb300-ea<br>Source code https://github.com/NVIDIA/TensorRT-LLM/tree/feat/gb300-ea<br>Early build with Beta quality - expect functional and performance issues since the support is in progress. |

### Table 9. 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** | NVIDIA DCGM |
| NVSSVT | **1.6** | NVOnline: **1108364** |
| nvdebug | **1.7.0** | NVOnline: **1109504** |
| nvfwupd | **2.0.7** | NVOnline: **1107320** |
| nvrastool | **1.5** | NVOnline: **1112947** |
| MSTflint | **4.34.0-2** | Github: MSTflint link |

# Improvements

This section provides information about the improvements in each release.

## Release 1.0.1

1. **Fixed an issue preventing modification of the Rsyslog TransportProtocol in the NVIDIA Switch BMC from default UDP to TCP or other protocols.** Transport protocol configuration is now supported.
2. **Resolved an authentication issue that led to a DoS-like state, blocking BMC access and affecting firmware updates and other functionality.**
3. **The method for identifying and managing GPUs that belong to the same compute tray has been updated.**
4. **IMEX issue in where outgoing gRPC connections were not reliably detected as lost, is improved by event handling and connection recovery, enabling nvidia-imex to promptly reestablish communication and resume processing when a disconnect occurs.**
5. **Excessive partition API requests can fill GFM logs.**

   Logging has been optimized to capture entries only during actual topology changes, ensuring the preservation of critical information.
6. **eth0/1 routing was not separated, causing incorrect routing when eth0 was disconnected**.

   Added a fix to separate the routing per interface (eth0/1).
7. **Resolved an intermittent issue where a deadlock in GFM would result in CUDA error "`All CUDA-capable devices are busy or unavailable`"**.
8. **Resolved an occasional failure with NVLS bind operation (`cuMulticastBindMem`)**
9. **FLR reset hangs during GPU engine transaction.**

   Fixed by leveraging the hardware engine reset state machine to reliably reset the engine.
10. **Fixed multicast reference count computation.**

    Fixed an issue that caused multicast reference count to be computed incorrectly.
11. **In rare instances, the PKEY request could be rejected by the subnet Manager.**

    Added a fix to remove the race condition and the PKEY database is always cleared before sending the response to the GFM.
12. **Partitions are deleted from the domain.**

    This update addresses the root cause of the GFM crash, ensuring stability and preventing recurrence of the issue.

13. **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. The removed GPU is set to a “reset required” state and can’t be used to run workloads until the GPU is reset.

14. **Resolved intermittent C2C link training failure during GPU reset.**

    During GPU reset, the Grace<->Blackwell GPU C2C link training 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, which ensures reliable C2C link training during reset.

15. **Resolved SGPI_D0 and SGPI_E0 status does not reflect MCIO J80.B12.**

    Added a fix to report the second presence pin, compared to using just one of the pins.

16. **This release includes PCIe link quality improvements that will resolve an occasional PCIe CTO (Completion Timeout) error that could result in a kernel crash.**

17. **Aligning the behavior of NMX-C API where duplicate partition creation returns RESOURCE_USED on GB300 vs. PARTITION_EXISTS in GB200.**

    With v1.0.1, GB300 will return NMX_ST_PARTITION_EXISTS (same as GB200).

18. **Updated NMX-C to return partitionId in UpdatePartitionResponse for RemoveGpusFromPartition ("User partition") and AddGpusToPartition ("User partition").**

    Earlier versions failed to return partitionId

19. **Resolved incorrect GPU THERM_WARN_INT error message in SEL logs while the GPU was operating still under normal temperature range.**

    Such an error was occasionally produced during stress test workloads. This was due to an incorrect temperature threshold in some of the code paths.

20. **Resolves an intermittent issue where a NCCL all-reduce throughput drops.** This occurs only when NVLink SHARP is enabled and control messages from GPU to GFM were dropped (for example, buffer contention), resulting in multicast setup failure. This version adds fixes across the GPU driver, Fabric Manager, and GPU firmware to improve the reliability of NVLS initialization.

# Release 1.0.0

1. **Fabric Manager improvements:**
    - GPU's fabric health summary is included in NVIDIA-smi/NVML fabric.health.summary on the compute tray.
    - Added a new row fabric.health.incorrect_configuration to include incorrect NVLink fabric configurations in nvidia-smi/NVML.
    - Assigned the GPU a unique clique ID when it's unhealthy or not included in the NVLink partition to block p2p traffic.
    - Trunk link admin state feature provides ability to manage trunk link states for maintenance purposes.
2. **SBIOS fixes**:
    - Resolved potential memory leaks in RAS_FW that can cause failures when attempting a large amount of GPU error injections
    - Resolved an interrupt handling issue in RAS_FW that could cause RAS_FW to miss an interrupt when handling GPU ECC errors
    - Resolved an issue in which the PSC ROM certificates occasionally have a signature encoding failure
3. **HMC improvements and fixes:**
    - Resolved an issue with missing alert events from secondary FPGA.
    - Resolved an issue with GB200 being displayed incorrectly on GB300 systems.
    - Resolved an issue with HMC not flipping i2c MUX to enable FPGA access to EROT SPI when FPGA_READY is asserted in certain corner cases.
    - Resolved an issue with incorrect handling of trailing "/" in Redfish URIs.
4. **BMC improvements and fixes:**
    - Resolved an issue with differentiating BMC reset reasons between normal and watchdog reset.
    - Resolved an issue with GB200 being displayed incorrectly on GB300 systems.
    - (5457362) Resolved an issue with missing BMC_ERoT_0 serial number when PowerPolicy was set to AlwaysOff.
    - Show all allowable ResetType values in BMC UI for System power operations.
    - Resolved an issue with /redfish/v1/TelemetryService/MetricReportDefinitions/{metric_report_definitions_instance} after POST /redfish/v1/TelemetryService/MetricReportDefinitions.
    - Resolved an issue with Virtual Media failing to mount ISO in the BMC web GUI.
5. **Resolved Xid143 Error 0x65 Due to FSP Boot Race Condition**
   A race condition in the FSP firmware occurred when out-of-band SPDM tasks from the Host Management Controller (HMC) accessed a staging buffer during the boot process. This conflict caused the staging buffer contents to fail authentication, leading to an Xid143. This issue is fixed in v1.0GA.
6. **XID119 GSP timeout during AC cycle**
   This release fixes an intermittent issue where a race condition within UEFI causes ATS programming to fail, as indicated by the following messages in dmesg:
   `NVRM: nvAssertFailedNoLog: Assertion failed: pGpu->getProperty(pGpu, PDB_PROP_GPU_ATS_SUPPORTED) @ mem_desc.c:2193`

7. **Resolve the issue for Unreachable Switch BMC After Power-Cycle Workflow Failure**
   This fix resolves an intermittent issue where an NVSwitch BMC became unreachable after a workflow attempted an auxiliary power cycle of compute and NVSwitch. The problem occurred when orchestration left the NVSwitch in an incomplete power state, making the BMC unresponsive. This update ensures NVSwitch power transitions are handled properly, maintaining BMC reachability after automated reboots.
8. **GPU Firmware Version Now Reported in Redfish FirmwareInventory**
   This release updates the SoftwareInventory "Version" field in Redfish UpdateService for GPUs, allowing for correct firmware identification.

# Known Issues

This section provides a list of the known issues.

1. **MC powercycle of compute and switch fail to bring up Fabric Manager.**

   In rare cases, following an unexpected power failure or a switch node crash, Fabric Manager may fail to start due to database (Fabric Manager resource information) corruption.

   **Workaround**: There is currently no method to recover the corrupted data. To restore functionality, perform an NMX-C reset to reinitialize the cluster and Fabric Manager.

2. **Segmentation faults occur when running some applications.**

   During process teardown, applications may encounter segmentation faults under some circumstances.

   **Workaround**: There is currently no workaround, and this issue will be fixed in TRD5.

3. **In rare instances, (< .01 %) BMC management interface may fail to train phy link.**

   In rare instances ( < 0.01%), the BMC management interface phy may not link correctly with a switch.

   **Workaround**: Reboot the BMC through the host via ipmi (ipmitool mc reset cold) or through Redfish via the BMC to Host USB Ethernet interface. ( curl -k -u "user:password" -X POST https://${bmcip}/redfish/v1/Managers/BMC_0/Actions/Manager.Reset -d '{"ResetType": "ForceRestart"}' )

4. **Unexpected data (PCB sensor) was detected "/redfish/v1/Chassis/HGX_Chassis_0/Sensors".**

   Renamed the PCB-Temp sensor to Exhaust-Temp sensor. While the PCB sensor no longer exists in the RF sensor list, the PCB temp events still exist and increment on exhaust temp event occurrence.

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

5. **When updating VBIOS, percent complete still shows 0, after the task is complete.**

   In some cases, the task service indicates a success condition while the progress property remains at ‘0’. In this case, ignore the progress property and monitor the TaskState and TaskStatus to determine update status.

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

6. **Grace needs to clear the dpc_trigger_status.**

   In hotplug scenarios that involve CX8 downstream devices that enable DPC, some DPC events that are triggered by the hotplug attempts are not correctly handled by RAS firmware. This causes the DPC trigger status to remain set after a hotplug DPC event, which results in the downXstream link remaining disabled.

   **Workaround**: Disable DPC for CX8 downstream links before you attempt to hotplug the devices. This issue will be fixed in the next release.

7. **NeighborMTUDiscards property is not defined in MetricReportDefinitions**

   NeighborMTUDiscards property is not defined in MetricReportDefinitions but exists in the Metric Reports.

   **Workaround**: This property will be removed in a future release. Property read will fail - do not use.

8. **When an incorrect static topology file(for example 2x36) is specified on a 1x72 topology, FM log displays duplicate control plane error states**

   When an incorrect static topology file(for example 2x36) is specified on a 1x72 topology, both CONFIG_ERROR_ADDITIONAL_CHASSIS_DETECTED and CONFIG_ERROR_MISSING_CHASSIS are set.

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

9. **GPU Status state missing from metric reports**

   Status/Status URI is missing from */redfish/v1/TelemetryService/MetricReports/HGX_ProcessorMetrics_0*

   **Workaround:** Status can be read from telemetry service processor metrics (/redfish/v1/Systems/HGX_Baseboard_0/Processors/GPU_{GpuId}#/Status/State). This issue will be fixed in a future release.

10. **Disable access link retraining**

    The GPU Access Link Retraining feature is disabled in this release. After upgrading, a compute tray reboot or GPU reset is required. With this release, a compute tray reboot or GPU reset is also required to bring GPU NVLink back online when a switch tray is rebooted or power-cycled for any reason. This feature will be re-enabled in a future NVOS release.

11. **Inconsistent GPU memory reported by NSM, nvidia-smi, and Redfish**

    Multiple management interfaces report different total GPU memory values: NSM Type 3 (0x0C) and nvidia-smi (-q -d MEMORY) return 284,208 MiB, while Redfish TotalMemorySizeMiB returns 285,324 MiB. This occurs when querying inventory or memory capacity from NSM, nvidia-smi, MODS, and Redfish.

    Reporting inconsistency may cause confusion in monitoring, inventory reconciliation, or capacity planning. There is no data-loss or security impact identified.

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

12. **BMC PCIe link reset causes SBIOS exception.**
    When the PCIe link between the Grace CPU and BMC is reset at runtime, the system might take a fatal exception, which generates a Fatal CPER event. This issue only impacts the NVIDIA reference design but might impact partners.
    **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.

13. **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.

14. Fan Control and Leak Detector user configured settings that are modified through Redfish PATCH API will reset to default after BMC firmware update.
    Some Fan Control and Leak Detector properties are configurable through the PATCH method on Redfish. For example, the user is able to modify whether the BMC will shutdown the Chassis when a leak is detected through this API:

    curl -s -k -u ${USER}:${PASSWORD} https://${BMCIP}/redfish/v1/Chassis/Chassis_0/Oem/Nvidia/Policies/LeakDetectionPolicy --request PATCH -d '{"PolicyEnabled":true}

    This setting will persist through BMC resets and tray power cycles, but will not currently survive a BMC firmware update.

    **Workaround**
    Check and reapply any desired settings after a BMC firmware update.

15. The Redfish Firmware Inventory API intermittent fails to fetch certain firmware endpoints
    The Redfish API to retrieve the current FirmwareInventory:

curl -s -k -u ${USER}:${PASSWORD}
https://${BMCIP}/redfish/v1/UpdateService/FirmwareInventory

Would be missing some firmware inventory endpoints on rare occasions. The cause has been identified and the fix will be included in a future BMC release.

**Workaround**

Re-run the same Redfish API.

## 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



NVIDIA Corporation | 2788 San Tomas Expressway, Santa Clara, CA 95051
http://www.nvidia.com