



NVIDIA DGX GB200 NVL72 2.0.0 Release Notes

***Release Notes for the 2.0.0 GA Multi-Node System Stack
Firmware and Software Components for NVIDIA DGX
GB200 NVL72***

NVIDIA Corporation

Oct 07, 2026

Contents

1	Document History	1
2	Overview	3
3	NVIDIA DGX GB200 NVL72 2 Features	5
4	Multi-Node System Software Stack Package Contents	7
4.1	Compute Tray Components	7
4.1.1	Host Software Components	7
4.1.2	BMC	8
4.1.3	HMC	8
4.2	GB200 Switch Tray Components	9
4.2.1	NVOS	9
4.2.2	BMC+FPGA+EROT	9
4.2.3	SBIOS+EROT	9
4.2.4	CPLD	10
4.3	PowerShelf FW	10
4.3.1	NVIDIA/gdrcopy	10
5	NVOnline IDs for System Software, Tools, and Drivers	11
5.1	Public Release Links Associated with this Release	11
5.2	cuDNN	11
5.3	NCCL	11
5.4	NVSHMEM	12
5.5	DLFW Containers	12
5.6	TRT-LLM	13
5.7	MFT/Tools Firmware Package Contents	13
5.8	Tested Compatibility	13
6	Improvements	15
6.1	Release 2.0.0 GA	15
6.1.1	BMC and HMC Improvements	24
6.1.1.1	bios-settings-mgr	24
6.1.1.2	openbmc	24
6.1.1.3	dbus-sensors	27
6.1.1.4	entity-manager	27
6.1.1.5	libmctp	27
6.1.1.6	linux	28
6.1.1.7	nsmd	28
6.1.1.8	nvidia-code-mgmt	28
6.1.1.9	bmcweb	28
6.1.1.10	nvidia-ipmi-oem	29
6.1.1.11	nvidia-power-manager	30
6.1.1.12	nv-shmem	30

6.1.1.13	phosphor-bmc-code-mgmt	30
6.1.1.14	phosphor-dbus-interfaces	30
6.1.1.15	phosphor-host-ipmid	30
6.1.1.16	phosphor-logging	30
6.1.1.17	phosphor-time-manager	30
6.1.1.18	phosphor-networkd	30
6.1.1.19	phosphor-state-manager	30
6.1.1.20	phosphor-user-manager	30
6.1.1.21	pldm	31
6.1.1.22	spdm	31
6.2	Release 1.3.10	31
7	Known Issues	33
7.1	DGX Issues	36
8	Notices	39
8.1	Notice	39
8.2	Trademarks	40
8.3	Copyright	40

Chapter 1. Document History

Version	Date	Description of Change
1	June 9, 2025	1.0 GA Release
2	June 30, 2025	1.0 GA Release - Switch ERoT Version typo fixed.
3	July 29, 2025	1.2.2 GA Release
4	September 30, 2025	1.3 GA Release
5	November 25, 2025	1.3.1 GA Release
6	January 9, 2026	1.3.2 GA Release
7	March 24, 2026	1.3.6 GA Release
8	July 10, 2026	1.3.10 GA Release
9	September 22, 2026	2.0.0 GA Release

Chapter 2. Overview

This document contains detailed information about the SW 2.0.0 GA release for the NVIDIA® DGX GB200 72x1 rack configuration 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 software 2.0.0 GA release is only for the target product DGX GB200 NVL72 and cannot be used for any other product.

Note

Action required: Upgrade for NVLink Recovery

In release 2.0.0, NVLink Recovery remains enabled by default. To ensure proper functionality and system stability, you must upgrade all components in the rack to release 1.3.2 or later, including compute nodes and NVSwitch.

Failure to upgrade the entire rack may lead to incompatibility issues and unexpected behavior during NVLink Recovery operations.

Note

Action required: Upgrade for TPM firmware 15.25

Infineon has published a firmware update for the SLB9672 TPM module used on DGX GB200 and DGX GB300 compute trays. The update addresses a known security vulnerability in the TPM firmware and is recommended by the NVIDIA SW and Security teams as a security update. The TPM ships disabled by default on NVIDIA DGX GB200 and DGX GB300 systems, so most deployed systems are not exposed to this vulnerability. The update is available if you have enabled and are using the TPM. Refer to the `README.txt` file included with the TPM firmware for update instructions.

Note

Action required: Upgrade for NVMe firmware

Optional Samsung firmware GDC7902Q for PM9A3 drives is available with this release for systems that use Samsung drives.

Refer to [Updating NVMe Device Firmware — NVIDIA DGX B300 Firmware Update Guide](#) for the NVMe Device Firmware update instructions.

Chapter 3. NVIDIA DGX GB200 NVL72 2 Features

This section lists the NVIDIA DGX GB200 NVL72 features in this release.

This release includes the following features:

Note

No new features were introduced in the 2.0.0 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 NVBandwidth.
- ▶ NCCL
 - ▶ NCCL: All2All, AllReduce, and Unicast.
 - ▶ NVSHMEM
- ▶ Math Libraries
 - ▶ cuBLAS (included in CUDA Toolkit 13.0 Update 2).
- ▶ DCGM
 - ▶ 4.6.0 release.
- ▶ NVSM
 - ▶ 25.09.07

Chapter 4. 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.

4.1. Compute Tray Components

Refer to the NVIDIA Multi-Node NVLink System Bring-up 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 DGX GB200 Release 2.0.0.

4.1.1. Host Software Components

Component	Version
GPU Driver	580.173.10
IMEX	580.173.10
CUDA Toolkit	13.0.3
MFT Tools	4.36.0-147
DOCA_Host	3.4.1-010000
NMX-M	85.1.1100
CX7	28.49.1118
BF3	3.4.1-11
Host OS	DGX OS 7.5.0

4.1.2. BMC

Version: 260710.1.0_custom

Bundle file: nvfw_DGX-GBX00_0023_260710.1.0_custom_prod-signed.fwpkg

Component	Version
EROT	01.04.0055.0000_n04
BMC	GB200Nv1-26.07-1

Component	Filename
CoRIM	DGX-GBX00_0023_260710.1.0_custom_prod-signed.corim
Recovery Firmware	nvfw_DGX-GBX00_0101_260710.1.0_custom_recovery_prod-signed.fwpkg

4.1.3. HMC

Version: 260710.1.0_custom

Bundle file: nvfw_HGX-GBX00_0023_260710.1.0_custom_prod-signed.fwpkg

Component	Version
CPLD	0.22
PS GPU	97.00.F8.00.1B
QS GPU	97.00.F8.00.19
EROT	01.04.0055.0000_n04
HMC	GB200Nv1-26.07-1
SBIOS	02.06.06
FPGA	1.66

Component	Filename
CoRIM	HGX-GBX00_0023_260710.1.0_custom_prod-signed.corim
CRT	nvfw_GB200-P4975_fwpkg_prod_cert.crt
Recovery Firmware	nvfw_HGX-GBX00_0023_260710.1.0_custom_recovery_prod-signed.fwpkg
PEM	nvfw_GB200-P4975_fwpkg_prod_pub.pem

4.2. GB200 Switch Tray Components

4.2.1. NVOS

Version: 25.02.4463

Component	Filename
CoRIM	GB200-P4978_0005_260611.1.0_prod-signed.corim

4.2.2. BMC+FPGA+EROT

Version: GB200-P4978_0004_260616.1.0

Bundle file: nvfw_GB200-P4978_0004_260616.1.0_recovery_prod-signed.fwpkg

Component	Version
EROT	01.04.0055.0000_n04
BMC	88.0002.1984
FPGA	0.24

Component	Filename
Recovery Firmware	nvfw_GB200-P4978_0004_260616.1.0_recovery_prod-signed.fwpkg
CoRIM	GB200-P4978_0004_260616.1.0_prod-signed.corim

4.2.3. SBIOS+EROT

Version: GB200-P4978_0006_260602.1.0

Bundle file: nvfw_GB200-P4978_0006_260602.1.0_prod-signed.fwpkg

Component	Version
EROT	01.04.0055.0000_n04
SBIOS	0ACTV_01.01.022

Component	Filename
CoRIM	GB200-P4978_0006_260602.1.0_prod-signed.corim
Recovery Firmware	nvfw_GB200-P4978_0006_260602.1.0_recovery_prod-signed.fwpkg

4.2.4. CPLD

Version: GB200-P4978_0007_260701.1.0

Bundle file: nvfw_GB200-P4978_0007_260701.1.0_prod-signed.fwpkg

Component	Version
CPLD	000370_REV0700000377_REV2100000373_REV1300000390_REV0400
FUI OPN	FUI000583

4.3. PowerShelf FW

Component	Version
Powershell PSC	0104
Powershell PMC	3.2.4

4.3.1. NVIDIA/gdrcopy

NVIDIA/gdrcopy is a fast GPU memory copy library based on NVIDIA GPUDirect RDMA technology. GDRCopy version 2.5 is publicly available.

- ▶ 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++

Chapter 5. NVOnline IDs for System Software, Tools, and Drivers

5.1. 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/TRD11)	580.173.02	13.0.3	CUDA driver 580.173.02 download
CUDA Toolkit 13.0.3		13.0.3	CUDA Toolkit 13.0.3 download
CUDA Toolkit 13.0.2 Download.		13.0.2	CUDA Toolkit 13.0.2 download
NVLink and CX7/BF3/DOCA	All Versions		NVIDIA Networking software documentation

5.2. cuDNN

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

Component	Version	Location
cuDNN	9.18.1	NVIDIA cuDNN release notes

5.3. NCCL

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

Component	Location
NCCL Home-page	NVIDIA Collective Communications Library (NCCL) Download Page
PIP Wheels	NVIDIA Collective Communication Library (NCCL) Runtime
GitHub release	NVIDIA NCCL GitHub release
DevZone Binaries	NVIDIA NCCL download page (current release) NVIDIA NCCL legacy download page (legacy releases)
Release notes	Overview and NCCL release notes Note: GB200 support is available from NCCL 2.27.5

5.4. NVSHMEM

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

Component	Location
NVSHMEM Homepage	NVIDIA NVSHMEM Homepage
Current Release	NVIDIA NVSHMEM Downloads
Legacy releases	NVIDIA NVSHMEM Archive of legacy releases
Documentation	Release Notes/Installation Guide: NVIDIA NVSHMEM Release Notes API Documentation: NVIDIA OpenSHMEM Library (NVSHMEM) Documentation Best Practice Guide: NVIDIA NVSHMEM Best Practice Guide

5.5. DLFW Containers

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

Component	Versions	Location
PyTorch	Latest tag from NGC catalog.	NGC catalog for PyTorch
TensorRT	Latest tag from NGC catalog.	NGC catalog for TensorRT
JAX	Latest tag from NGC catalog.	NGC catalog for JAX

5.6. TRT-LLM

Component	Version	Location
TRT-LLM	Current version	TensorRT-LLM
vLLM	25.10	N/A
FlashInfer	0.2.6	N/A

5.7. 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	4.5.0	NVIDIA DCGM
nvdebug	2.0.0	Installation Guide — NVIDIA Debug Tool for Datacenter Products
nvfwupd	2.0.8	nvfwupd tool version v2.0.8
MSTflint	4.34.1-3	GitHub: MSTflint link

5.8. Tested Compatibility

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

Upgrade / Downgrade
1.3.10 <-> 2.0.0

Chapter 6. Improvements

This section describes the improvements in each release.

6.1. Release 2.0.0 GA

1. This release includes the following features:

- ▶ NVLink Down reset improvements
- ▶ Enhanced Fabric state reporting for GPUs with misprogrammed CBC EEPROM
- ▶ Enhanced GFM API error reporting
- ▶ Single Switch Tray recovery with adaptive bandwidth

These features require the following minimum versions:

- ▶ GPU driver 580 TRD6 or later
- ▶ 2.0 recommended SW/FW stack

2. Refer to the NVOS Release Notes in the NVOS NVOnline posting for additional information about the improvements in this release.
3. This release includes security updates. Refer to the Partner Security Bulletin: NVIDIA Blackwell Data Center Products - June 2026.
4. This driver release resolves two conditions where the GPU driver could report XID 109 spuriously:

On B300 and later GPUs, a driver issue could occur when multiple processes were using the GPU while ECC counters were being monitored through NVML or `nvidia-smi`. In rare instances, one of the processes could stop making progress, causing the driver to report XID 109 and terminate the affected process.

On Hopper and later GPUs, when a process encounters a GPU error, such as a `device-side assert()`, the process is terminated as expected. However, if GPU core dumps are enabled, `CUDA_WAITS_ON_EXCEPTION` is set, or a debugger is attached, process teardown could time out and cause the driver to report XID 109. Other processes resume after the teardown notification.

5. NCCL Hangs with Coherent GPU Memory Mode

Resolved an issue where enabling `NVreg_CoherentGPUMemoryMode=driver` on GB200 caused NCCL applications to hang. A CDMM mode detection bug caused the driver to incorrectly detect the mode as NUMA, which affected dma-buf mapping because this mode is handled differently. The driver now detects the mode correctly. (Fixed in R580 TRD7.)

6. Traffic after compute tray swaps.

Resolved an issue where traffic could fail after compute tray swaps across locations within a partition. This fix supports the documented customer workflow for compute tray swaps within a partition.

7. NVLink retraining during GPU probes.

Resolved an issue where NVLink retraining could conflict with GPU probe request and response handling within a partition. This conflict could disrupt workloads when retraining and GPU probe operations occurred at the same time.

8. Multicast workloads fail after a factory reset.

Resolved an issue where multicast workloads could fail after a factory reset because GPU handles were not recognized correctly.

9. NVLink range logging.

Corrected Fabric Manager log output so NVLink link ranges are reported accurately for easier diagnosis.

10. SError After CBB Timeout

Resolved an issue where handling a CBB timeout to PCIe configuration space could incorrectly trigger an Arm processor SError and kernel panic.

11. ATF Workaround for Arm Neoverse V2 Erratum 2743011

Resolved a logic error in the ATF workaround for Arm Neoverse V2 erratum 2743011. The error could prevent the workaround from being applied even when the boot message indicated that it was active.

12. SKU Authentication Bypass in PLDM firmware update.

Resolved a security vulnerability in the PLDM firmware update state machine where a command sequence could allow an adversary to flash firmware intended for a different hardware SKU without passing authentication checks. The fix updates the pending trigger table when firmware data is written to flash, preventing the authentication bypass.

13. Backup-Partition Wear-Out from PLDM Cancel Flow.

Resolved an issue where repeated PLDM UpdateComponent -> write -> CancelUpdateComponent cycles could leave partial flash data and consume erase cycles on the same backup-partition sectors. The fix adds a rate-limit check for UpdateComponent.

14. EROt hangs from VDM Multipacket buffer overflow.

Resolved an issue where oversized VDM middle fragments could overflow the composer buffer, causing the overflow path to return true with the output packet pointer left as nullptr. Because the VDM driver path lacked the null pointer check used by the PLDM and SPDm paths, this could result in a null pointer dereference and an approximately 4 GB memcpy, causing the EROt to hang.

15. Potential code execution in EROt DOT provisioning path.

The fix applies the correctly bounded logic from `auth_metadata_buffer()` to the parallel `update_metadata_sig()` DOT path.

16. InfoROM Read-Only Strap Policy Bypass in PLDM.

Resolved an issue where `pldmfw_cmd_pass_component_table()` recorded a component ID in the `pass_comp_id()` table before evaluating the strap gate. The fix updates `pass_comp_id()` only after the component ID check completes.

17. EROt denial of service in Non-MCTP PLDM adapter.

Resolved an issue in the non-MCTP PLDM adapter. The fix eliminates the vector that could disrupt root-of-trust attestation and firmware update services.

18. Fixed an issue with NVBMC GPU*_PWRO sensors reporting 'na' when the GPU power is > 1300 Watts.
19. SPDMM CPU attestation fails when MCTP response exceeds one packet. Resolved an issue where SPDMM CPU attestation could fail with a certificate-length mismatch when an MCTP response was one byte longer than a single-packet message. In this case, EROt marked the first packet of the multi-packet response with EOM set to true. Multi-packet response length handling now correctly accounts for this case.
20. Glacier EROt firmware detects SPI flash devices without hardware write protection.
Added support in Glacier EROt firmware to detect and report Serial Peripheral Interface (SPI) flash devices that do not support hardware write protection through the SRWD/SRP bit.
21. FirmwareState and BackgroundCopyStatus consistency in DOT volatile mode.
Resolved an issue in DOT Volatile Mode where, after a background copy completed, BackgroundCopyStatus could report Completed while the active slot FirmwareState still reported WriteInProgress and the inactive slot incorrectly reported Activated. NSM now prioritizes the AP boot-complete slot when determining the active slot, keeping active and inactive slot reporting aligned.
22. Added support for CEC1736 N05 CPN.
Added support for the CEC1736 N05 CPN, including a dedicated S8 key set for the N05 part.
23. FPGA-EROt I2C driver strength reduced to prevent SDA undershoot.
Resolved an issue on HGX-B200 systems where the FPGA-EROt I2C target driver strength could cause the SDA line at the customer BMC chip to undershoot below the -0.5 V GND limit during read operations. No functional failures were observed in NVIDIA lab measurements, but customer testing reproduced the signal-integrity violation. EROt firmware now supports a new VDM command to reduce the I2C driver output strength for the FPGA-EROt at address 0x52.
24. Improved EROt firmware upgrade logs for revoked security keys.
Improved EROt firmware upgrade logging when an upgrade fails with Security keys revoked. The logs now provide clearer information for key revocation and rollback-protection cases.
25. EC firmware debug key revocation uses updated NSM cache.
Resolved an issue where EC firmware debug key revocation could fail because NSM used stale cached key index and permission values after boot. NSM could return 0x0 instead of the production-configuration values key_idx=1 and key_revocation_info=0x01, causing firmware to determine that no revocation was needed. NSM now populates the cache from flash before processing key-revocation commands.
26. NVLOOM workloads fail after removing an unallocated GPU from a partition.
Resolved an issue where NVLOOM workloads could fail after approximately 30 minutes when two 1-GPU VMs were running in a 3-GPU partition and the third, unallocated GPU was removed from the partition. The failure could trigger a kgsprcAndNotifyAllChannels_IMPL() crash and log an Xid 94 error in one VM. This issue is fixed in the 6.17.0-1015-nvidia-64k kernel and later.
27. GPU reset operations no longer require power cycle after NVSwitch tray reboot.
Resolved an issue where GPU reset operations could fail with kernel Xid errors, such as GSP time-out, when a subset of NVSwitch trays rebooted while an NVLink SHARP workload was running. In some cases, recovery required a rack-level AC or DC power cycle.

28. NVSwitch ERoT forced recovery works on Juliet switch board.

Resolved an issue where forcing a functional NVSwitch ERoT into recovery mode and then force updating it did not work on the Juliet switch board. The forced recovery and update flow now supports this case. Standard ERoT functionality, failed-ERoT recovery, and normal firmware update mechanisms were not affected.

29. RDMA NCCL failures resolved after RDM-GFM communication loss.

Resolved an issue where, in rare cases, communication loss between RDM and GFM could cause a memory leak in NVLSM and lead to RDMA NCCL failures. NVLSM now handles this communication-loss scenario without requiring nmx-c processes to be restarted.

30. Micron NVMe drives remain available behind CX8 with PCIe max payload configuration.

Resolved an issue where Micron M.2 and E1.S NVMe drives could become unavailable on PCIe when the system used the NVIDIA reference SBIOS and all installed drives were Micron. The SBIOS now programs the PCIe root ports with the proper PCIe Max Payload setting. GB200 systems with CX7 were not affected.

31. OS image PCIe errors and system reboots are resolved.

Resolved an issue where a PCIe surprise link down between the Grace root port and NIC could make the NIC unreachable, cause cascading service failures in Kubernetes, Cilium, Node Cache, and Node Exporter which leads to a kernel panic after function-level reset (FLR) and PCIe retrain recovery attempts failed. This issue is fixed in DOCA 3.4.0. Also resolved an issue where, during kdump kernel invocation after the kernel panic, the Grace SMMU could enter a global abort state, terminate in-flight transactions, emit Poison TLPs, and cause Completion Timeouts. This issue is fixed in the 7.0.0-1015-nvidia-64k kernel and later.

32. HMC memory no longer fills during update startup.

Resolved an issue where repeated calls to the temp sensor shell script could leak small amounts of memory and eventually deplete HMC memory when starting an update. This could trigger the out-of-memory (OOM) killer and cause the HMC to crash or behave unexpectedly. The memory leak in the temp sensor shell script is fixed.

33. Deepseek 1N job causing Xid 48 and fatal node crash.

Deepseek workload fails causing an OS crash with RM driver 580.159.01. This is not an issue with RM driver 580.126.20 TRD7 or RM driver 580.126.16 TRD6.

Workaround: Perform a GPU reset or driver unload/reload in between workloads.

34. DPC limit from BMC incorrectly changed GPU power under power brake.

Applying a DPC limit from the BMC could alter GPU power consumption while the power brake was engaged. Limits are now deferred so they take effect only after the power brake disengages, as intended.

35. One or more GPU IRoT(s) not found.

On GPU I3C tasks, the implementation had an issue where unknown tasks would lead to failure. This issue is resolved and implemented in the I3C recovery flow.

36. Unsupported Target GPU T.Limit sensor disabled.

The Target GPU T.Limit sensor is not supported but could still report an invalid threshold of 0xFF. Recent firmware disables this sensor so the invalid value is no longer returned.

37. Clear Max Observed Power NSM API ignored aggregate clear (0xFF).

The Clear Max Observed Power NSM API (message type 3, command code 5) did not accept the aggregate argument 0xFF to clear all max power readouts. Added support for that will clear the aggregate argument driver branch 595.28 and newer.

38. XID 150/154 immediately after GPU reset (MSE communication initialization).

During reboot stress testing (22+ hours), XID 150 and 154 could rarely appear immediately after a GPU reset. A race condition in the MSE communications path could leave a message request to MSE in an “empty but incomplete” state. Without recovery for that state, the communications flow did not reinitialize properly and MSE became unable to process new requests. Recovery handling for this state has been added.

39. MCU firmware update can fail on a CX8 port when MCU with TRNG issue fails to boot.

On CX8, the MCU did not always update to the correct firmware on one port. When the MCU had a TRNG-related issue, it could fail to boot, which caused the firmware update to fail. The update path was made more resilient by adding multiple retries so the MCU can complete the update when boot is delayed or initially unsuccessful.

40. Enhance Logging for C2C training and memory training failure.

This release includes logging enhancements to allow more debug information in ROT log.

41. Redfish telemetry on HMC interrupted by GPU OOB NSM timeouts.

Resolved an intermittent timeout between the HMC and GPU OOB firmware that disrupted Redfish telemetry when NSM replies were delayed. The underlying timing issue is fixed in VBIOS firmware, restoring stable telemetry collection.

42. CPU cores may get stuck in firmware leading to a CPU softlock under the operating system.

Resolved issues in ATF, RAS, and SatMC firmware that could lead to a CPU core getting stuck in the firmware, causing a CPU softlock under the operating system.

43. Hard lockup due to uncorrectable error consumed by GSP.

Fixed an issue in RAS_FW that caused a CPU core to be stuck in an infinite loop.

44. Null-valued GPU sensors.

Under some conditions, internal processing of requests can take more time than expected. During this time, another request may be received from the BMC. In this scenario, OOBHUB will respond to the wrong request if a response for the first request is sent. Implemented a mechanism to drop the request if it exceeds the timeout threshold.

45. Intermittent GPU boot failure.

Fixed a GPU boot failure that occurred when FLR was issued while NVLink was training.

46. “Route Unhealthy” state remains True after reroute().

Repeated NVLink in-band messages could exhaust credits in the driver because of a credit-handling bug. In that situation, the route could stay in “Route Unhealthy” even after a reroute. The driver’s credit handling has been corrected so that credits are not exhausted in this scenario and the route returns to a healthy state after reroute.

47. NMX-C unhealthy, GFM gRPC times out.

Resolved an issue where multiple threads in the GFM had a circular dependency on locks, causing a deadlock. GFM could not make progress in that state, which led to NMX-C appearing unhealthy and GFM gRPC calls timing out. The deadlock has been removed so GFM can run normally and gRPC requests complete.

48. Node swaps no longer require fabric restart (GUID-based partitions).

For default GUID-based partitions, tray replacement now correctly removes the old tray's GUIDs and adds the new tray's GUIDs, so a fabric restart is no longer required after a node swap. For user-defined GUID-based partitions, the partition create and add-GPU APIs now enforce the NVLink domain GPU limit so partition size cannot exceed the domain limit.

49. GPU containment with no related data, XIDs, C2C, or NVLink errors.

Improved the logging of the GPU Containment CPER to include additional register dumps that are useful when analyzing failures.

50. C2C bandwidth out of range, unexpected interrupt and buffer mismatch.

Fixed an issue in which UEFI-reserved memory may be included in the OS's available memory map when EGM is enabled.

51. CPU core soft lockup.

Fixed multiple issues in SBIOS firmware that could cause a core to get stuck in the firmware for extended periods of time, leading to the OS reporting soft lockups.

52. NVRASTool v1.2.1 OOB full error injection – 7 fail, 5 pass when run individually.

HMC now acknowledges GPIO alerts from both primary and secondary FPGA. Resolved a synchronization mismatch between the FPGA and HMC firmware: the FPGA drove separate GPIO alerts for primary and secondary controllers, while the HMC only monitored an aggregated signal from the primary FPGA. Alerts from the secondary FPGA were therefore not acknowledged by the HMC. HMC firmware has been updated to monitor GPIO alerts from both the primary and secondary FPGA so all controller alerts are acknowledged.

53. Failed to get HGX_CPU_ sensors.

Resolved this issue by creating associations to CPU inventory path in entity-manager without relying on SmbiosMDR service.

54. Redfish reporting “disabled” status for multiple CPU ports.

Fixed an issue in SatMC that causes PCIe ports to incorrectly report “disabled” in Redfish reporting.

55. GPU-PGOOD is missing in the metric report.

Added the GPU's status.state property for GPU-PGOOD, which was previously missing in the GPU ProcessorMetrics report.

56. GPU-Status CPER logged at runtime with Redfish fatal severity.

Changed the GPU-Status CPER severity from recoverable to informational to ensure that this informational event is handled correctly by some operating systems.

57. MemorySpareChannelPresence reported as false when information is unavailable.

When MemorySpareChannel information could not be read (for example when SatMC is down), the HMC set the Presence (state) to false by default, which incorrectly suggested that no spare channel was present. The default is now unknown in the backend and null in Redfish when the information is unavailable, so the state correctly reflects that the value is unknown rather than false.

58. No P2P after reboot following Xid119 on FW1.2.1 (tray swap/replacement).

When a tray was replaced or swapped, the internally used systemImageGuid stayed with the physical location instead of moving with the tray. That led to duplicate nodeIds, corrupted remaps, and failed P2P or workload failures (including no P2P after reboot following XID 119).

on FW 1.2.1). Fabric Manager now uses a software-based approach called “gpu set” to identify GPUs that belong to a single tray. This correctly handles trays that have been moved to a different location and new GPU trays added in a previous location, eliminating the duplicate nodeld and remap corruption issues and restoring P2P after tray swap or replacement.

59. Node swaps no longer require fabric restart.

Resolved an issue in how GPUs on the same tray are identified. In one code path, GPUs could use an uninitialized nodeld, which incorrectly grouped them with uninitialized GPUs from other trays and led to REMAP and P2P failures, often requiring a fabric restart after node swaps. Tray identification is now correct in that path, so REMAP and P2P no longer fail for this reason and node swaps do not require a fabric restart.

60. RemoveGpusFromPartition no longer fails when partition is already unhealthy.

Previously, if an NVLink partition’s routing was unhealthy, the NMX Controller failed all partition-modification API calls (add or remove GPU), so no changes could be made until the partition was healthy again. The controller now fails and reverts only when the API call itself would make the partition unhealthy. RemoveGpusFromPartition (and add-GPU calls) can therefore succeed when the partition is already unhealthy, allowing administrators to modify the partition in that state.

61. RFE support for reduced bandwidth when a switch tray is offline.

When a switch tray goes offline (due to a fault or admin action), the fabric can now continue operating using the remaining switches at reduced bandwidth. In earlier releases, the NVLink domain became non-operational in this case. This feature currently supports a single switch tray being absent, provided that tray is not the primary switch tray running the NMX Controller.

62. RFE: GPU NVLink traffic when moving GPUs across partitions without GPU reset.

This enhancement to Multi-Node NVLink partitioning allows administrators to move GPUs between NVLink partitions for multi-tenancy without requiring a GPU reset or GPU resident service restart to enable peer-to-peer (P2P) traffic over NVLink.

The behavior is controlled by the new Fabric Manager option `MNNVL_NVLINK_PARTITION_MODE`, which supports two modes:

- ▶ Multi-tenancy NVLink Partition Mode (Default value = 1): A GPU reset is required when a GPU is moved between partitions. In this mode, `nvidia-smi -q` shows GPU Recover.

Action: Drain and Reset until the reset is performed.

- ▶ Trusted NVLink Partition Mode (Value = 2): GPUs can be moved between partitions and use P2P traffic immediately, without a reset or service restart. In this mode, `nvidia-smi -q` shows GPU Recover.

Action: None during the move.

Changes to this setting take effect after an NMX-C restart.

63. CreatePartition with duplicate GPUs did not return PARTITION_EXISTS.

Fixed an issue where Fabric Manager did not return PARTITION_EXISTS when creating a partition that would reuse GPUs from an existing partition. The cause was default resiliency mode not being set on creation, so the existence check compared unset vs set values and failed to match. Default resiliency mode is now set correctly on partition creation, so the correct error is returned.

64. Reroute returns generic GRPC error (NMX_ST_GENERIC_ERROR).

Fixed an issue where the partition reroute API returned NMX_ST_GENERIC_ERROR for most failures. New error codes have been added so that specific failure scenarios return the appropriate code, improving diagnostics and error handling.

65. Fabric Manager uplink calculation for partially populated topologies.

Resolved an issue where Fabric Manager (FM) could miscalculate required uplinks in partially populated topologies, so probe could fail even when enough links were present. The uplink calculation now correctly accounts for available links in these topologies, allowing probe to succeed when the required number of uplinks is satisfied.

66. GFM underflow when allocating multicast resources to a partition.

Fixed an issue where GFM could miscalculate remaining multicast teams when allocating MC resources to a partition due to an unsigned-integer underflow in the availability check; remaining resources are now computed correctly.

67. Multiple control plane status description string.

The control plane status was incorrectly set to `CONFIG_ERROR_ADDITIONAL_CHASSIS_DETECTED` when only one chassis was discovered, even though the topology file indicated two chassis. The control plane status description is now set according to the actual discovered chassis count.

68. GFM log flooding due to constant topology updates.

Admin state logging is optimized so that admin states are logged only when they change. This prevents GFM log flooding caused by constant topology updates.

69. NV-SW Fabric Manager not reflecting correct GPU health/state.

GPU health states were incorrectly reported as healthy through the NMX-C API when the compute node was powered off or rebooted. The fix checks for the presence of GPU nodes before computing their health, so the Fabric Manager now reflects the correct state.

70. RFE - Enhance GFM APIs to GPU/Switch/Node NVLink states and transaction IDs.

NMX-C APIs now return structured, detailed error information for failed fabric and partition operations instead of only a generic error code. NMX-C APIs can now be used as the main interface for fabric troubleshooting and no longer need to manually correlate cryptic GFM logs.

71. Restricting PMU reset controls.

Allowing external PMU reset at elevated priority could, in some cases, expose a timeout bug. This change disallows external reset requests at elevated priority to eliminate that risk.

72. Adding a GPU to a location-based partition failed.

A bug in SM caused adding a GPU to a location-based partition to fail. A bug in GFM cleanup logic prevented subsequent add operations from succeeding. This fix ensures GFM cleans up correctly after failures so that later add attempts succeed. The SM bug is fixed in SM 1.3.1.

73. RemoveGpusFromPartition failed with `NMX_ST_GENERIC_ERROR`.

Removed the constraint that prevented removing GPUs with active multicast teams from a partition. Those GPUs can now be removed and are marked as "Needs Reset."

74. Partition remained unhealthy after NMX-C restart trunk link routing was lost.

Previously, when a partition became unhealthy because of a trunk link routing issue, restarting NMX-C or rebooting the switch that hosts NMX-C did not trigger a reroute. Users had to manually start the reroute to restore the partition. With this fix, restarting NMX-C or rebooting the hosting switch automatically triggers a reroute of trunk links and restores the partition to a healthy state when no trunk link issues persist.

75. `GPU_THERM_WARN_INT-ERROR` during CSST at safe temperatures.

Adjusted temperature-sensor warning thresholds so `THERM_WARN_INT` is not reported when the GPU is within safe operating limits.

76. Support for MSE LFM to notify GFM alive state.

In multi-node systems, GPUs previously timed out on NVLink detection if compute trays started before GFM/NMX-C, requiring reboots or GPU resets when NMX-C came online causing down-time.

MSE 13.1 introduces a heartbeat-based notification in LFM to detect GFM connectivity and immediately alert RM to trigger fabric probing, replacing timeout-based retries. This enables GPUs to establish NVLink connectivity with late-starting NMX-C without resets or reboots, eliminating operational disruption.

77. The system cannot boot up.

Fixed an issue that may prevent boot when power the brake is asserted.

78. [5394666] FWERROR reported in PEI.

Fixed an issue that would cause all FWERROR CPERs to report PEI mode.

79. RoT HGX_ERoT_BMC_0 (HMC) certificate extraction issues.

Resolved this issue by adding retry for certificate request when truncated response is received in spdman connection handler.

80. HMC reported invalid data (SlotIndex and SlotLocation).

Resolved this issue by adding retry to get UUID command and enhanced MCTP rediscovery.

81. HMC BMCweb hardcoded check for update header in MultipartHeader.

Changed the content-disposition field of multipart firmware update to non case sensitive. Prior to this change, the field "Content-Disposition" required exact casing to get recognized.

82. Host based copies regression with 1.1.

The nvbandwidth tool now supports testing with Huge Pages for host memory. Users can enable this feature by adding the `-useHugePages` flag to the command line. This enhancement will be available starting with the nvbandwidth v0.9 release.

83. Graceful handling of empty sensor list.

When the host is off, CPUs are expected to have no sensors or metrics. Previously, this returned an internal error instead of an empty list. Added a fix to remove the error response and now it returns an empty list when no sensors are found.

84. Telemetry missing for the power brake.

Added additional throttle reason updates for hardware slowdown, thermal slowdown, power brake slowdown, and sync boost in the `NsmProcessorThrottleReason::updateReading` function. This improves the processor's performance interface by providing more detailed throttle reasons.

85. HIP-1 down training to 1x1 after the DC cycle with the BM BIOS flavor.

Introduced a SatMC PLDM effector to allow out-of-band software to trigger a warm reset of the CPU.

86. Reporting NVLink BW performance variations across all nodes.

The `cudaLaunchAttributeNVLinkUtilCentricScheduling` launch attribute, accessible via `cudaLaunchKernelEx()` to indicate that a kernel should prioritize maximizing NVLink bandwidth utilization. When applied to communication-intensive kernels, such as those used by NCCL, CUDA automatically adjusts scheduling to better saturate NVLink bandwidth across modern GPU architectures. This enhancement helps mitigate performance variability in collective communication workloads observed on newer hardware platforms.

87. System occasionally does not report BERT error after injecting a CSF UE.

Added support for reading records by index and consuming multiple RAS log records. This enables more reliable error reporting by capturing multiple records that may have been generated between collections.

88. MSE implementation to support node description setting in GPU.

Previously, the NMX-T Prometheus interface returned hardcoded values (“GB100 Nvidia Technologies”) for `node_description` and `host_name`, preventing the ability to identify individual GB100 nodes and correlate metrics in OCI monitoring. MSE 14.14 now supports GET/SET_HOST_NAME messages to set hostnames in NetIR. Monitoring systems now display correct hostnames.

6.1.1. BMC and HMC Improvements

This release also includes the following BMC and HMC improvements:

6.1.1.1 bios-settings-mgr

- ▶ Corrected `ScureCurrentBoot` to `SecureCurrentBoot` in Redfish and the BIOS manager.

6.1.1.2 openbmc

- ▶ (HMC) Removed manufacturing script-exec and parse-yaml recipes.
- ▶ (HMC) Updated mbedtls to 3.6 to fix CVE vulnerabilities in SPDH.
- ▶ (HMC) Fixed HMC Redfish Status.State incorrectly reporting Starting after boot.
- ▶ (BMC) Flushed journal logs to eMMC on aux power cycle and after firmware update to prevent log loss.
- ▶ (HMC) Fixed memory leak in hmc-temp-sensors daemon.
- ▶ Disabled ssh pubkey auth to prevent ssh login after debug token removal (CVE-2025-23337).
- ▶ Removed slp-lite service (CVE-2024-41660).
- ▶ Added maximum recovery attempts limit for BMC/HMC filesystem recovery.
- ▶ Improved FPGA1 FW recovery process.
- ▶ If host power is off during PCIe switch update, ignored update failure and returned success to make it consistent with GPU FW update behavior.
- ▶ (DGXOPENBMC) Reapplied built-in user settings and group permissions from ROFS after firmware update.
- ▶ Fixed issues with PCIe switch on/off power flow.
- ▶ Updated HMC to Kernel 6.12.41.
- ▶ Added new metrics for GPU, CPU, HMC, and FPGA.
- ▶ Differentiated BMC reset types: Enabled kernel watchdog pretimeout panic to distinguish between normal reset, kernel panic, and watchdog timeout.
- ▶ Fixed Redfish event SSE filtering when message ID is missing.
- ▶ Fixed randomly generated password for bootstrap account not always complying with password complexity requirements.

- ▶ Selected correct Metric Report Definition Metrics depending on CX7 or CX8 configuration.
- ▶ Added missing virtual I2C buses to FRU device service blocklist.
- ▶ Added support for GB200 NVL4.
- ▶ Fixed missing support for alerts from the second FPGA.
- ▶ Fixed CX7/CX8/BF3 FW RelatedItems in Redfish.
- ▶ Fixed redundant CPU id issue in CPU Metric Report Definition.
- ▶ Removed service account from BMC and disabled -S option for SCP.
- ▶ Removed service account from BMC.
- ▶ Disabled SCP -S option.
- ▶ Added support for GB200 NVL4.
- ▶ Added support for additional CX8 variants.
- ▶ Enabled kernel watchdog timeout detection SEL logging.
- ▶ Fixed issues with firewall rules preventing access to SSH ports 2201-2203 to connect to the Grace CPU UARTs.
- ▶ Enabled IPMI command filtering based on host state. Added allow list for CPU/SBIOS with ProvisionedHostAllowlist configuration.
- ▶ Updated 1RU exhaust and pdb inlet thermal setpoints based on latest thermal recommendations.
- ▶ Added support for additional GB200 variant.
- ▶ Fixed issues with missing ipmi oem command support with GB200/GB300 DGX builds.
- ▶ Fixed CPU max power setting in IPMI/Redfish causing incorrect sensor value reporting.
- ▶ Updated BF3 temperature sensor reporting to use PLDM and instead of i2c which fixed occasional crash of hwmontempsensor service.
- ▶ Fixed issues with PDB FRU being occasionally lost during host reboot.
- ▶ Fixed issues with psusensor service slowly leaking thread descriptors
- ▶ Enabled ERoT log collection via I2C fallback if PLDM methods fail.
- ▶ Fixed issue with built-in admin account becoming unusable if deletion of the account is attempted. Now throw 403 errors on deletion attempts.
- ▶ Disabled deprecated faultlog in bmcweb since this was moved to EventLog.
- ▶ Fixed issue with LastResetTime not always being synced with chassis power state.
- ▶ Fixed issue with missing shutdown options displayed in the WebUI.
- ▶ SSIF interface reliability enhancements.
- ▶ Enabled OOB clearing UEFI variables without reboot.
- ▶ Enhanced WebUI SOL to include host power state indicator.
- ▶ Enhanced WebUI to filter non-present E1.S drives as status absent by default instead of status critical.
- ▶ Enhanced error handling, logging, and reporting when the emmc log partition is corrupted or unavailable.
- ▶ Modified CPLD FW update to not report error if host main power is off.

- ▶ Fixed kernel watchdog timer reset reporting and logging.
- ▶ Fixed issue with OpenOCD service not always auto-disabling when JTAG interface is removed.
- ▶ Enabled 10-second fan controller watchdog timer to ramp fans to 100% if BMC fails to pet the watchdog.
- ▶ Fixed kernel crash issue with attempted access to eeprom devices when in incorrect host power state.
- ▶ Fixed kernel APSEED driver issues with unhandled i2c IRQs.
- ▶ Fixed issue with CPER event logging.
- ▶ Fixed Redfish BMC and HMC product name reporting.
- ▶ Fixed missing CPU FW RelatedItem in Redfish.
- ▶ Fixed Redfish FW Inventory HGX_FW_CPLD RelatedItem pointing to incorrect Chassis item.
- ▶ Fixed Redfish related items for firmware inventory entries.
- ▶ Patched BMC kernel to include SSIF workaround for data mismatch handling to prevent i2c-tegra driver crashes in host OS.
- ▶ Fixed various missing RelatedItems in Redfish FW Inventory.
- ▶ Fixed WebUI FW Update pages to allow FW update in any host power state with clearer warnings about what may not be supported in different power states.
- ▶ Fixed BMC ERoT FW version occasionally missing after DC power cycle.
- ▶ Added support for additional factory diags tests.
- ▶ Added support for NVL4.
- ▶ Fixed issue with BMC ERoT serial number missing after HMC reboot.
- ▶ Fixed issue with missing CPU sensors in certain corner cases.
- ▶ Fixed occasional phosphor-create-chassis-poweron-log service failure during soft reboots.
- ▶ Fixed issue with user creation through Redfish when password complexity requirements were not met.
- ▶ Added EffectiveError to ProcessorPortMetrics MRD.
- ▶ Fixed issue with ProcessorPortMetrics containing removed properties.
- ▶ Fixed missing HGX_MetricProperties for GPU.
- ▶ Added support for additional GB300 variant.
- ▶ Fixed issue with ERoT SPI access post FPGA0 recovery if AC power cycle was not performed post recovery.
- ▶ Fixed issue with missing GPU PN and SN reporting.
- ▶ Fixed intermittent issue with pldmd and fansensor coredumps.
- ▶ Added logic to persist host side MAC address of BMC-to-host USB Ethernet adapter link.
- ▶ Added additional CX8 Power Telemetry sensors.
- ▶ Fixed issue with missing GPU thermal sensor readings.
- ▶ Fixed issue with certain NVMe drive sensors being lost after DC cycle.
- ▶ Updated thermal setpoints for ConnectX ports.

- ▶ Fixed intermittent issue with false 0 RPM critical FAN readings events.
- ▶ Fixed intermittent issue with invalid fan critical threshold events being reported after cycling AUX power.
- ▶ Fixed missing Redfish CX8 power telemetry metrics.
- ▶ Updated thermal recommendations for PDB sensors.
- ▶ Fixed intermittent issue with failing to log BMC reset reason after BMC reboot.
- ▶ Report NVLink LED state in Redfish.
- ▶ Fixed intermittent issue with chassis-poweron-log service failing to start after BMC reboot.
- ▶ Enabled OOB boot variable clear (that is, CMOS clear).
- ▶ Aspeed RTC only re-initializes on cold boot so warm reboot preserves time and HMC stays synchronized with the BMC.
- ▶ Enabled network adapter device protection across D-Bus, NSM, and Redfish.
- ▶ Added configured warning/critical/shutdown temperature limits for Quanta PDB HSC sensors.
- ▶ Included FPGA I3C APB register dump output in HMC FPGA diagnostics.
- ▶ Added GPU PCIe port inventory in entity-manager so Redfish processor GPU ports expose PCIe_0 as well as NVLink.

6.1.1.3 dbus-sensors

- ▶ Fixed intermittent issue with fansensor service coredumping during host power cycle.
- ▶ Fixed intermittent issue with nvmesensor service coredumping after bmc reboot.
- ▶ Fixed intermittent issue with nvmesensor falsely reporting -1 temperature.
- ▶ (HMC) Fixed intermittent issue with lpcsnop and nvidia-fpga1-ready services failing to start.

6.1.1.4 entity-manager

- ▶ Fixed intermittent issue with duplicate path handling leading to coredumps.

6.1.1.5 libmctp

- ▶ (BMC and HMC) Fixed OOB heap read in libmctp Get UUID response handlers that could leak adjacent heap memory via Redfish.
- ▶ (BMC and HMC) Fixed OOB heap reads in multiple MCTP control response handlers.
- ▶ (DGXOPENBMC) Enabled systemd journal APIs for MCTP logging.
- ▶ Fixed issue with duplicate MCTP control packet handling causing CX8 discovery to fail.
- ▶ Fixed occasional issue with CX7/CX8 Temp reporting as 4000/8000C due to rare error handling issue with MCTP over I2C that could lead to duplicate messages being transmitted over the bus. This could occasionally lead to PLDM message response confusion.
- ▶ Capped consecutive SPI AP receive timeouts so mctp-demux can return to its main loop and satisfy systemd watchdog.
- ▶ Removed Redfish logging for restart notify failure for the ERoT in certain spurious cases.

6.1.1.6 linux

- ▶ (HMC) Moved to Linux kernel 6.12.73.
- ▶ (BMC) Hold BMC kernel at 6.12.9 revision due to kernel instability. Kernel will be updated after the root cause is found.
- ▶ Aspeed RTC only re-initializes on cold boot so warm reboot preserves time and HMC stays synchronized with the BMC.

6.1.1.7 nsmd

- ▶ (BMC and HMC) Fixed buffer overflow vulnerabilities in libnsm response decoders.
- ▶ (HMC) Fixed invalid (NaN) sensor readings reported as valid OC data by setting the telemetry stale bit.
- ▶ (HMC) Fixed GPU assembly FRU fields (part number, serial number) returning blank after a DC power cycle.
- ▶ (HMC) Fixed PowerSmoothing AdminOverride mask issue where it was not fully cleared after ApplyAdminOverrides.
- ▶ Various NSMD bugfixes and performance improvements
- ▶ Enabled various missing telemetry items.
- ▶ (HMC) Fixed missing power brake telemetry.
- ▶ (HMC) Fixed issue with GPU utilization metrics.
- ▶ (HMC) DGXOPENBMC Reduced High Latency in Redfish Workload Power Profile Management.
- ▶ (HMC) Reduced High Latency in Redfish Workload Power Profile Management.

6.1.1.8 nvidia-code-mgmt

- ▶ (HMC) Fixed heap OOB read in firmware image signature verifier when parsing an undersized firmware component.
- ▶ (HMC) Fixed heap out-of-bounds read in firmware image hash verification caused by a 32-bit integer overflow.
- ▶ Fixed issue with fw-status crashing sometimes when changing power states.
- ▶ (BMC) Fixed issue when using targeted FW update. PCIeSwitch would always be updated if contained in the bundle even if not part of the targeted update.
- ▶ Fixed issue when using targeted FW update. PCIeSwitch would always be updated if contained in the bundle even if not part of the targeted update.
- ▶ (HMC) Fixed issue with BMC firmware update in ERoT-less systems.

6.1.1.9 bmcweb

- ▶ (BMC) Fixed Redfish SSE subscriptions not starting the BMC/HMC event aggregation listener, causing an SSE event stream to receive no events until a separate EventDestination subscription was created.
- ▶ (BMC) Fixed incorrect task-progress status check for the SPI erase operation.
- ▶ (BMC and HMC) Fixed Redfish Service Validator conformance errors in OEM ActionInfo and schema resources.

- ▶ (BMC) Added missing odata.type to SPDM signed measurements attestation response.
- ▶ (BMC) Fixed Redfish log entry collections to sort numerically by ID, correcting broken log pagination.
- ▶ (BMC and HMC) Fixed CPER event severity code handling in the SSE and event log paths.
- ▶ Enabled IPMI command filtering based on host state. Added allow list for CPU/SBIOS with ProvisionedHostAllowlist configuration.
- ▶ Fixed issue with Empty Part Number reporting in Redfish.
- ▶ Fixed issue with incorrect handling of trailing "/" in Redfish URIs.
- ▶ Fixed issue with secondary SPI missing in inventory.
- ▶ Corrected ScurCurrentBoot to SecureCurrentBoot in Redfish and the BIOS manager.
- ▶ Fixed missing power cycle option reporting in uxPowerResetActionInfo.
- ▶ (BMC and HMC) Fixed issue with POST action timeout and Redfish aggregation.
- ▶ (BMC and HMC) Improved memory usage during firmware update.
- ▶ Fixed issue with occasional missing events with Redfish Event subscriptions.
- ▶ Fixed issue that could cause phosphor-networkd to crash and be exploited to cause a DoS of BMC services.
- ▶ Fixed issue with missing CPU sensors in certain corner cases.
- ▶ Fixed issue with incorrect RF return codes when NVMe drive cannot be accessed.
- ▶ Added new LsbIBGUID property.
- ▶ Fixed issue with RF aggregation when using VariableSpiRead.
- ▶ Fixed Redfish conformance issue with SSD temperature reporting when value is unavailable.
- ▶ Fixed issue with CPU sensor reporting when the host is powered off.
- ▶ Fixed memory leak in certain cases with Redfish aggregation.
- ▶ (BMC and HMC) Fixed display of EDPp persistence field value.
- ▶ (BMC and HMC) Fixed issue which contains links and HGX_ProcessorModule.
- ▶ (HMC) Report NVLink LED state in Redfish.
- ▶ Fixed Redfish location of LeakDetector policy ReactionDelaySeconds PATCH operations to comply with schema.
- ▶ (HMC) Fixed missing Redfish OriginOfCondition field in certain event log entries.
- ▶ (BMC) Fixed issue with Redfish server not closing inactive TCP sessions.
- ▶ Enable network adapter device protection across D-Bus, NSM, and Redfish.
- ▶ Fixed intermittent corrupted MetricReportDefinition property names under aggregation (URL fragment use-after-free when rebuilding URIs).

6.1.1.10 nvidia-ipmi-oem

- ▶ Fixed randomly generated password for bootstrap account not always complying with password complexity requirements.

6.1.1.11 nvidia-power-manager

- ▶ (DGXOPENBMC) If host power is off during PCIe switch update, ignored update failure and returned success to make it consistent with GPU FW update behavior.

6.1.1.12 nv-shmem

- ▶ Enabled various missing telemetry items.
- ▶ Fixed missing PCIe gen6 type in pcie-type map.

6.1.1.13 phosphor-bmc-code-mgmt

- ▶ Added logic to persist host side MAC address of BMC-to-host USB Ethernet adapter link.

6.1.1.14 phosphor-dbus-interfaces

- ▶ Fixed issue that could cause phosphor-networkd to crash and be exploited to cause a DoS of BMC services.
- ▶ Fixed issue with missing CPU sensors in certain corner cases.
- ▶ (HMC) Report NVLink LED state in Redfish.
- ▶ Enabled network adapter device protection across D-Bus, NSM, and Redfish.

6.1.1.15 phosphor-host-ipmid

- ▶ Enabled IPMI command filtering based on host state.
- ▶ Fixed issue with BIOS reset not being reliably cleared.

6.1.1.16 phosphor-logging

- ▶ Fixed issue when handling corrupted data logging files.

6.1.1.17 phosphor-time-manager

- ▶ Changed EROt timesync failure event log to informational rather than critical since it does not impact operation.

6.1.1.18 phosphor-networkd

- ▶ Fixed issue that could cause phosphor-networkd to crash and be exploited to cause a DoS of BMC services.

6.1.1.19 phosphor-state-manager

- ▶ (DGXOPENBMC) Validated cpu boot done signal.

6.1.1.20 phosphor-user-manager

- ▶ Fixed issue with built-in admin account becoming unusable if deletion of the account is attempted. Now throw 403 errors on deletion attempts.

6.1.1.21 pldm

- ▶ (DGXOPENBMC) Fixed issue with logging FW_UPDATE_BLOCKED messages in Redfish.
- ▶ (HMC) Fixed issue with GPU FW update not being skipped without error when the host is powered off.

6.1.1.22 spdm

- ▶ (HMC) Updated mbedtls to 3.6 to fix CVE vulnerabilities in SPDM.
- ▶ (BMC and HMC) Fixed OOB read in SPDM NEGOTIATE_ALGORITHMS request decoder by bounding checking algorithm counts.
- ▶ (BMC and HMC) Fixed the issue where SPDM requests failed promptly with a clear error when a device was unreachable instead of hanging for minutes.
- ▶ (BMC and HMC) Fixed OOB heap read in SPDM CERTIFICATE response decoder.
- ▶ (BMC and HMC) Updated mbedtls crypto library to version 3.6 for the SPDM daemon.

6.2. Release 1.3.10

1. Refer to the NVOS Release Notes in the NVOS NVOnline posting for additional information about the improvements in this release.
2. This release includes security updates. Refer to the Partner Security Bulletin: NVIDIA Blackwell Data Center Products - June 2026.
3. GPU recovery after failed AddGpusToPartition call.
Resolved an issue where a GPU could fail to rejoin the NVLink fabric after a failed AddGpusToPartition API call during an add or recovery flow.
4. Multicast team teardown during HA events.
Resolved an issue where multicast persistence handling could prematurely tear down multicast teams during high availability (HA) resiliency events.
5. OpenBMC HMC memory leak when starting an update.
Resolved an issue where the hardware temperature sensor service could leak memory on the OpenBMC HMC when an update was started.
6. SError after HSM_FABRIC fault.
Resolved an issue that could cause an SError to occur after an HSM_FABRIC error. The fault is now handled correctly without raising a spurious SError.
7. Disabled HBM SBE reporting by default.
HBM SBEs are now disabled by default to avoid duplicate reporting as they are already reported via XIDs. Otherwise, use the BCT knob to enable this reporting to continue double reporting of HBM SBEs.
8. Neoverse V2 Erratum 2618597.
Included Arm's suggested errata workaround for 2618597.
9. Resolved an issue with SDEI error reporting.
Save NS context before existing RAS handling of SDEI events to avoid TLS pointer corruption.

10. XID 150/154 immediately after GPU reset (MSE communication initialization).

During reboot stress testing (22+ hours), XID 150 and 154 could rarely appear immediately after a GPU reset. A race condition in the MSE communications path could leave a message request to MSE in an “empty but incomplete” state. Without recovery for that state, the communications flow did not reinitialize properly and MSE became unable to process new requests. Recovery handling for this state has been added.

11. Some GPUs may fail to ramp to full TGP when power smoothing ramp-up is enabled, resulting in reduced compute performance.

Implemented a VBIOS fix to update the ramp-up logic to allow the ramp-up.

12. Enhance Logging for C2C training and memory training failure.

This release includes logging enhancements to allow more debug information in ROT log.

13. Intermittent GPU boot failure.

Fixed a GPU boot failure that occurred when FLR was issued while NVLink was training.

14. Micron NVMe drives remain available behind CX8 with PCIe max payload configuration.

Resolved an issue where Micron M.2 and E1.S NVMe drives could become unavailable on PCIe when the system used the NVIDIA reference SBIOS and all installed drives were Micron. The SBIOS now programs the PCIe root ports with the proper PCIe Max Payload setting. GB200 systems with CX7 were not affected.

15. NVSwitch CPU reboots with MGMT board VR failure.

Added an improvement to avoid an NVSwitch CPU reboot due to a BMC forced/unexpected reboot.

16. XID 145 due to incorrect routing configured by NVLSM (36x2 only).

Resolved an issue where NVLink Subnet Manager (NVLSM) configured incorrect routing in switches in one rack of 36x2 towards GPU ALID when GPU is belonging to another rack of 36x2.

17. Power smoothing minimum power floor increase.

This enhancement to power smoothing now allows the Minimum Power Floor (MPF) value to be configured to 95% (previously 90%).

18. Workarounds for Arm Neoverse-V2 erratum 3841324 and 3442699.

Implemented suggested workarounds for two Arm erratum that previously caused incorrect instructions from being executed.

19. Some GPUs may fail to ramp to full TGP when power smoothing ramp-up is enabled, resulting in reduced compute performance.

Implemented a VBIOS fix to update the ramp-up logic to allow the ramp-up.

Chapter 7. Known Issues

This section lists the known issues in this release.

1. Trusted NVLink partition is not supported in this release.
2. Adaptive Bandwidth Mode and NVLink retraining require a TRD11 driver or later.
3. Compute node firmware update from an older version (any version before 0.9) will require an intermediate upgrade step before upgrading to the latest version of firmware.

There is a GPU security change that requires customers to flash from v0.6 -> v0.8 -> v0.9 -> 1.x

Workaround: As an interim step, go to the v0.8 GPU firmware.

4. Firmware Update task completed 100% but says critical.

The full log causes the task status update to fail.

Workaround: The fw update was successful, no workaround is needed. To confirm that the status changes to successful, clear the eMMC and run it again:

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

5. 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 deasserted).

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

6. The HMC_BMC_0 Status is stuck in a Starting state.

The HMC MCTP service signaling of "start" was not being triggered properly to move the task to the completed state.

Workaround: Check the HMC_READY GPI and FirmwareInventory in Redfish for another way to determine whether the HMC is ready.

7. ERoTs are not responding to MCTP.

A secure boot task cancels background copy. When the background copy was cancelled, the ERoT was responding to MCTP with errors. This issue with ERoT v1.4.7 was addressed in ERoT v1.4.8 and in the GB200 NVL72 GB200NVL72_P4975_CX7_Release_0.9.00 full release.

Workaround: Flash back to ERoT .262 or move forward to ERoT 1.4.8.

Multiple AX cycles might be required to get ERoT "unstuck" and allow a firmware update. To avoid race conditions, disable BMC Auto power on.

8. FPGA ERoT 1 is missing in the DOT Redfish APIs.

All Incoming PG548 (Bianca) boards from NVIDIA L6 have one sockmet (1P) SBIOS. During L10 manufacturing, if a system boots with two 1P SBIOS, module 0 boots but module 1 does not, which triggers the HMC to only build Redfish for one of the two modules.

Workaround: Disable BMC Auto power on and do not let the system boot until a proper 2S SBIOS is flashed. After a system enters this failed state, the HMC should not show module 1 in Redfish until after an HMC factory reset and AUX power cycle is performed.

9. HMC Telemetry Missing in Redfish.

This issue impacts some partner BMCs due to an I2C contention issue that is reading the HMC FRU on I2C bus 3. The HMC FRU is on a shared I2C bus between the BMC and the HMC and can lead to a multi-master collision.

Workaround: Before accessing the FRU and checking Redfish to the HMC, ensure that the BMC is properly checking for HMC_READY.

10. SBIOS flash interruption during a Grace power cycle.

During an SBIOS flash, if Grace is reset, the ERoT will abort the flash and give SPI priority to boot the CPU. This causes Boot Chan 0 to become corrupted. SBIOS boots from Boot Chain 1, and a background copy will recover the redundancy.

Workaround: Complete one of the following tasks:

- Ensure that the SBIOS flash is finished before resetting Grace.
- The BMC can hold Grace in reset by asserting `SYS\_RST` until the flash is done.

11. FPGA ERoT 1 is missing in the DOT Redfish APIs.

All Incoming PG548 (Bianca) boards from NVIDIA L6 have one socket (1P) SBIOS. During L10 manufacturing, if a system boots with two 1P SBIOS, module 0 boots but module 1 does not, this triggers the HMC to only build Redfish for one of the two modules.

Workaround: Disable the BMC Auto power on and do not allow the system boot until a proper 2S SBIOS is flashed. After a system gets in this failed state, the HMC will not show "module 1" in Redfish until after an HMC Factory Reset and an AUX power cycle is performed.

12. The SBIOS will not boot on after an L2 reset.

If Boot Chain is corrupt or will not boot, the Grace/EROT will try 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.

13. Parallel GPU resets led to unexpected system crashes a single digit percentage of time.

A GPIO-based GPU reset depends on a four-way handshake that is based on the GPIO pins being in sync. However, there are corner cases where this handshake can get out of sync, which might cause RAS_FW to make an incorrect assumption about the reset status of the GPU, and cause a fault.

Workaround: If RAS_FW detects that the GPU init status is questionable, RAS_FW will request an additional reset from users.

14. cuSOLVERMp 0.9.0 workloads may encounter GPU errors.

Workloads that use cuSOLVERMp 0.9.0 may encounter a GPU error during execution. In affected cases, driver logs may report XID 31 followed by XID 109.

Workaround: Update the workload to use the latest cuSOLVERMp.

15. Intermittent CPU-only HPL performance variability.

On some systems, certain 2.0 release candidate firmware stacks may show intermittent CPU-only HPL performance instability, including higher run-to-run variation and lower sustained FP64 results than expected.

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

16. Intermittent topology-related failure during GPU partition recovery after reset.

During GPU partition recovery after reset, an intermittent race between Fabric Manager and the switch manager can cause premature reset handling and a topology-related failure before recovery fully completes.

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

17. Intermittent NVLink SHARP multicast registration failure on GB300 multi-node systems.

On multi-node systems, NVLink SHARP multicast setup may intermittently fail after abnormal partition deletion, causing multicast memory registration to fail with an invalid-argument error and leading to `ncc1CommWindowRegister` failure.

Workaround: Avoid using a partition ID that was used previously.

18. Performance degradation for 256 MB NCCL AllReduce messages with MNNVL disabled.

On RoCE configurations running NCCL 2.30.7 with MNNVL disabled, the NCCL tuner may select NVLSTree/Simple instead of Ring/Simple for 256 MB AllReduce messages. This selection can cause lower AllReduce performance for affected workloads.

Workaround: Set `NCCL_ALGO=RING` and `NCCL_PROTO=SIMPLE` for affected workloads.

19. Multi-property PATCH to Redfish AccountService returns HTTP 400 for one out-of-range value.

When a single PATCH request updates multiple Redfish AccountService properties and one value is below the enforced security minimum, such as `AccountLockoutThreshold` or `AccountLockoutDuration`, the valid properties are still applied. However, the response returns HTTP 400 instead of the HTTP 200 response with a per-property error required by the Redfish specification.

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

20. FPGA may permanently assert power brake after a secondary module overt fault.

If a thermal OVERT condition occurs on the secondary module, the FPGA may incorrectly assert power brake permanently. This condition can reduce system performance on subsequent power-ups.

Workaround: AC power cycle the compute tray to clear the condition.

21. BMC ERoT serial number may fail to populate in Redfish (< 0.5%).

Workaround: This does not cause any operational issues. Perform a graceful restart of the BMC to correct the serial number display.

22. GPU re-introduced the following correction of location programming. Link status is not updated.

Fabric Manager detects GPUs that have invalid or duplicate locations, usually indicating an EEPROM mis-programming and does not initialize these GPUs. In the duplicated location scenario, both GPUs are introduced into the cluster when the information is corrected. The port states may not be updated in this scenario for these GPUs being introduced, causing messaging and/or NVLS failures.

Workaround: Perform a GPU reset of the GPUs being introduced to correct this error.

23. On certain 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`). The behavior is caused by a race condition in a driver, which can be triggered when the internal InfiniBand debug utility `iblinkinfo` runs on NVOS. NVIDIA plans to include a fix in a future NVOS release.

24. `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 on the `FWID` value contained in the DICE certificate.

In the following example:

- GPU CoMID value for measurement block 4 is:

```
"sha-384;g0Q20Htc9drxCgG6pdBzpQ6SHuQY9YVsV78zy0P9GG3Vg6N6TyUnmIqz/
zchkTkM"
```

which is

```
"80E436D07B5CF5DAF10A01BAA5D073A50E921EE418F5856C57BF33C8E3FD186DD583A37A4F2527988A"
```

in hexadecimal.

- The DICE `FWID[0]` value is:

```
"BB5FEC93F7EC3E89204A21F4B1AF6B941E92A2B5550DA06009C06455439CF16C6EEAC3234DCBB3A4A"
```

The example Python code is:

```
import hashlib
fwid = bytes.fromhex(
    'BB5FEC93F7EC3E89204A21F4B1AF6B941E92A2B5550DA06009C06455439CF16C6EEAC3234DCBB3A4A'
)
msr = bytearray(48)
msr_buf = bytearray([0x0, 0x0, 0x11, 0x0])
msr_buf.extend(fwid)
msr_buf.extend(msr)
msr = hashlib.sha384(msr_buf).digest()
print(msr.hex())
```

Output:

```
80e436d07b5cf5daf10a01baa5d073a50e921ee418f5856c57bf33c8e3fd186dd583a37a4f2527988a
```

7.1. DGX Issues

1. 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, you can set `PowerRestorePolicy` to either `AlwaysOn` or `AlwaysOff`. Otherwise, you can send an additional power-on or power-off command after the power is restored to achieve the desired host power state.

2. Leak detection voltage sensors do not show the 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
```

3. No mechanism to disable the 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.

4. 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.

5. 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.

6. 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.

7. 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.

8. Leak detector voltage values read through the compute tray BMC via IPMI are clamped to 0.165 and 1.815.

For leak presence/absence, use the standard leak detection interfaces under `/redfish/v1/Chassis/Chassis_0/ThermalSubsystem/LeakDetection/LeakDetectors`. Do not use raw sensor voltage as the primary leak indicator.

Workaround: To read the correct voltage when the sensor is outside that range, use Redfish (for example, `/redfish/v1/Chassis/Chassis_0/Sensors`) instead of IPMI.

Chapter 8. Notices

8.1. 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.

8.2. 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.

8.3. Copyright

© 2026 NVIDIA CORPORATION & AFFILIATES. All rights reserved.

Copyright

©2022-2026, NVIDIA Corporation