



NVIDIA Data Center GPU Driver version 535.129.03 (Linux)/ 537.70 (Windows)

Release Notes

Table of Contents

Chapter 1. Version Highlights.....	1
1.1. Software Versions.....	1
1.2. Fixed Issues.....	1
1.3. Known Issues.....	2
Chapter 2. Virtualization.....	17
Chapter 3. Hardware and Software Support.....	19

Chapter 1. Version Highlights

This section provides highlights of the NVIDIA Data Center GPU R535 Driver (version 535.129.03 Linux and 537.70 Windows).

For changes related to the 535 release of the NVIDIA display driver, review the file "NVIDIA_Changelog" available in the .run installer packages.

- ▶ Linux driver release date: 10/31/2023
- ▶ Windows driver release date: 10/31/2023

1.1. Software Versions

For this release, the software versions are as follows:

- ▶ CUDA Toolkit 12: 12.2.2
Note that starting with CUDA 11, individual components of the toolkit are versioned independently. For a full list of the individual versioned components (for example, nvcc, CUDA libraries, and so on), see the [CUDA Toolkit Release Notes](#).
- ▶ NVIDIA Data Center GPU Driver: 535.129.03 (Linux) / 537.70 (Windows)
- ▶ Fabric Manager: 535.129.03 (Use `nv-fabricmanager -v`)
- ▶ NVFlash: 5.791

For more information on getting started with the NVIDIA Fabric Manager on NVSwitch-based systems (for example, NVIDIA HGX A100), refer to the [Fabric Manager User Guide](#).

1.2. Fixed Issues

- ▶ 4309921 – Due to a race condition in the NVIDIA GPU driver, NVLink-connected systems, such as DGX H100, may intermittently boot with restricted P2P capabilities (e.g. a specific pair of GPUs reporting P2P as not supported). This issue has been present in all earlier 535 drivers and is now fixed with R535TeslaRD5.
- ▶ 4288650 – Fixed a bug where the driver sometimes reports "Trying to free already-free IRQ" during a failed initialization. The driver would also not clean up interrupt resources completely during exit.

- ▶ 4278367 – Fix for false GSP-RM RPC timeouts (XID 119) in case the CPU thread processing the RPC gets de-scheduled for a long time.
- ▶ 4272659 – A design defect has been identified and mitigated in the GPU kernel-mode driver, related to the GPUDirect RDMA support in MLNX_OFED and some Ubuntu kernels, commonly referred to as the [PeerDirect](#) technology, i.e. the one using the [peer-memory kernel patch](#). In specific scenarios, for example involving the cleanup after killing of a multi-process application, this issue may lead to use-after-free and potentially to kernel memory corruption.
- ▶ 4260289 - Fixed an issue that would prevent the UVM kernel module building correctly against the CentOS Stream 8 kernel and possibly other kernel versions due to an undefined reference to the `MPOL_PREFERRED_MANY` symbol.
- ▶ 4259080 - Fixed a regression that caused CUDA initialization failure on Grace Hopper with newer Linux kernel versions due to UVM incorrectly detecting ATS support was not available when it was.
- ▶ 4228552 – Asymmetric resets of GPUs and NVSwitches caused NVLinks to fault. Checks were added to reset links in fault and re-launch ALL.
- ▶ 4169030 – Asymmetric resets of GPUs and NVSwitches caused NVLinks to go into contain state. Checks were added to reset links in contain state and re-launch ALL.
- ▶ 4165142 – Fixed an issue arising from a race condition between internal GPU microcontrollers in fetching GPU memory temperature which causes the current HBM temperature to return an invalid value which is above the memory maximum operating temperature. This produced a negative T.Limit temperature reading.

1.3. Known Issues

General

- ▶ "Change ECC State" and "Enable Error Correction Code" do not change synchronously when ECC state changes. 3838953
- ▶ 4254491 – CUDA kernels that use the sparsity feature of tensor cores through the [mma.sp](#) PTX instruction on Hopper architecture GPUs may intermittently experience silent data corruption resulting in incorrect results. NVIDIA libraries currently do not provide access to tensor cores with sparsity so only kernels directly developed using the `mma.sp` PTX instruction are impacted. This issue will be fixed in an upcoming release.
- ▶ MPAM Support

MPAM is not supported in this release. To enable MPAM support in the future, changes to the SBIOS are in progress. To use the MPAM capabilities at that time, you need an MPAM-enabled kernel.
- ▶ `cpufreq-info` is not reporting correct core frequency on Grace. 4253549

The CPU frequency reported by the Linux kernel might vary significantly from the actual value.

Workaround There is currently no workaround, and this issue will be resolved in a future kernel version.

- ▶ Unable to allocate full GPU memory. 4278263

On CPU+GPU coherent memory platforms, such as Grace Hopper, GPU memory is added to the Linux kernel as a NUMA node and is managed by the kernel. The kernel minimum watermark threshold might prevent GPU applications from allocating the full GPU memory.

Workaround

To allocate full GPU memory, users can manually tune `/proc/sys/vm/min_free_kbytes`.

- ▶ CPU TempAvg Low event log Threshold message is wrong. 4299244

When the temperature goes below the low warning/critical threshold, there is a Redfish event log with the threshold value. The expected threshold value should be the low warning/critical threshold, but the high warning/critical threshold value is selected instead.

As a result, users cannot get the correct threshold value from the Redfish low warning/critical Redfish event.

Workaround When you inspect a low warning/critical Redfish event, get the correct low threshold value from the sensor URI (`/redfish/v1/Chassis/CPU_0/Sensors/ProcessorModule_0_CPU_0_TempAvg_0`).

- ▶ [IPMI] Duplicate sensor name. 3999921

Sensor names are unique but, as per the current specification, IPMI has a 16-character for the name field. Sensors that have the same initial 16-characters appear as duplicate sensors in the SDR.

Workaround

There is currently no workaround, and this issue will be resolved in a future release of the BMC. Until then, you can safely ignore the duplicate sensor.

- ▶ A driver library mismatch error after a successful installation and subsequent reboot causes GPU programs to not start. 4302438

GPU programs work immediately after installing the driver, but fail to load after rebooting, and messages like the following appear in the kernel log:

```
API mismatch: the client has the version xxx.xx but this kernel module has the version yyyy.yy. Please make sure that this kernel module and all NVIDIA driver components have the same version.
```

This is most likely because kernel modules from a previous driver version are in the `initramfs`.

Workaround

Rebuild the initramfs and reboot the system. If a reboot is not immediately possible, you can temporarily work around the issue by unloading and reloading the NVIDIA kernel modules.

- ▶ Grace is doing a seemingly unnecessary munmap of the untouched host memory when CUDA context is active. 4052424

Customers who use older stable/longterm trees should update to the latest upstream subversion, and customers who are using another tree should cherry-pick the subversion from one of the trees (versions linux-6.5.y, linux-6.1.y, or linux-5.15.y).

Workaround

There is currently no workaround.

- ▶ On GH200, a `pyt_bart` inference perf regression occurs after a driver update. 4308011

On GH200 systems, copies between pageable memory and the device can be offloaded to copy engines. While this provides many performance and functional benefits, a subset of copies that involve addresses that are not resident in the main memory were observed to be slow. To mitigate this performance bottleneck, a performance optimization that actively fetches the non-resident pages into memory before the copy operation starts has been introduced.

The pre-fetching operation involves some overhead, so users can tweak the page read latency, and then the paging operation starts. Users can set the threshold value for page read operations beyond which CUDA will actively try to page in non-resident pages. The default value is 0.008 milliseconds.

To reduce the number of page-in ops, users can set the value to 0.01 or 0.02 milliseconds in the `CUDA_BUFFER_PAGE_IN_THRESHOLD_MS=` environment variable.

- ▶ `ipmitool mc` information does not show the firmware versions. 4299371

The NVIDIA OpenBMC Reference solution does not provide firmware inventory versions through IPMI.

You can find the firmware inventory versions in the Redfish interface.

- ▶ System not powering back on when performing power cycle from BMC. 4305610

When performing a power cycle command from the BMC, such as IPMI Chassis Power Cycle, the system will shut down correctly, but it might not automatically restart.

Workaround

Manually restart the system after the shutdown. This issue will be addressed in a future NVIDIA OpenBMC Reference Firmware release.

- ▶ Some errors are set after a clean boot. 4312911

The Correctable Error (the `CorrErr` bit) and Unsupported Request (the `UnsupReq` bit) in the Device Status register (part of the PCI Express Capabilities set of registers)

and the Advisory Non-Fatal Error (the AdvNonFatalErr bit) in the Correctable Error Status register of the Advanced Error Status (AER) set of registers are getting set because the enumeration process occurs in the Linux kernel. During this process, when there is a read for a non-existent device downstream of the root port, the kernel receives a response as an Unsupported Request, so the UnsuppReq bit is set.

Since Grace is configured to treat UnsuppReq responses as Advisory Non-Fatal Error, and the AdvNonFatalErr bit is set, there will not be an Uncorrectable Error. Instead, a Correctable Error is raised, and the orrErr bit is set. Grace is configured to have a software leaky bucket for all the correctable errors, so unless the number of correctable errors observed is greater than 32 in one minute, the errors are not reported to the OS. As a result, these errors are set after a clean boot.

Workaround Currently, there is no workaround available.

- ▶ GPU firmware update fails with an "Update.1.0.VerificationFailed" message during upgrade-downgrade testing. 4332232

If the GPU accesses its SPI flash at approximately the same time the ERoT attempts to verify the GPU firmware image, the ERoT might time out. The PLDM firmware update might fail with VerifyResult equal to 0x99, which is reported in the Redfish task status as Update.1.0.VerificationFailed.

Workaround

1. Ensure a proper firmware update sequence and status checking is implemented for each iteration of upgrade-downgrade testing.
 2. After each firmware update, issue A/C power cycle for the firmware update to take effect
 3. After the A/C power cycle, wait for the host to be ready before verifying the firmware inventory.
 4. After the A/C power cycle, wait for background copy completion before initiating the next firmware update.
- ▶ Support for 4k page size in nv-p2p.

The NVIDIA driver's kernel mode [GPUDirect RDMA APIs](#) that are used for Peer-direct support in [MLNX_OFED](#) and [GPUDirect Storage](#) are not supported on GH200 platforms when used with Linux kernels that are configured with the 4K page size. These APIs are not functional and might lead to a kernel memory corruption.

Users are strongly encouraged to move their software stack to the [dma-buf APIs](#), which require the [open-source GPU driver](#), Linux kernel 5.12 or later (if upstream kernel GPUDirect RDMA support for RDMA networking is required), and NVIDIA Turing™ + GPU. Since the dma-buf APIs work correctly on 4K page kernels, using those APIs is the ideal mitigation for this issue.

- ▶ Reduced fan speed does not display with correct severity in the event log. 4241946

When running the NVIDIA BMC reference code, the severity of a fan under the critical threshold is not correct. It should be Critical, but it is reported as OK in Redfish.

Workaround

There is no workaround unless you modify the BMC reference source code.

- ▶ Graphics functionalities, such as EGL, GLX, and Vulkan, are currently not supported under the 4K OS page size. 4333780

Workaround

We recommend that you use a 64K page size.

- ▶ The peak message rate varies greatly. 4329895

In some circumstances, PCIe strict ordering might not work well in Grace Hopper. Users might experience low message rates when the PCIe write traffic targeting GPU memory is high.

Workaround

You can use PCIe relaxed ordering, for example, by using `ibv_reg_mr` with the `IBV_ACCESS_RELAXED_ORDERING` flag.

- ▶ NaNs reported in longevity tests when ECC error injection and video error injection are enabled. – 4051897

Although rare, during longevity and stress testing, some sensors might randomly display NaN readings. However, this issue typically occurs only once for a few readings, and subsequent readings during the sensor polling loop return to normal. So far, we have not encountered any cases where a sensor displayed consecutive NaN readings.

To handle occasional NaN readings, the host BMC or any client software that consumes sensor information should implement a retry mechanism so that it is resilient.

- ▶ `Set_declare SMBPBI Master caps` will return success. – 3977576

The undocumented SMBPBI opcode 0x8 does not return `ERR_OPCODE` in conformance with the SMBPBI specification. This opcode is not supported on the Hopper HGX 8-GPU baseboard and should not be used.

- ▶ Fail to pass Redfish Protocol Validator (1.1.6). – 4109466

The Redfish validator test case `POST /redfish/v1/AccountService/Accounts` fails because of a Redfish validator error, but there is no impact on the system functionality.

Work around

No workaround is required.

Refer to the <https://github.com/DMTF/Redfish-Protocol-Validator/issues/57> validator issue for more information.

- ▶ Instead of a failed action, a success action is displayed when a AccountLockoutDuration PATCH request is sent in AccountService. – 4109596

When a PATCH request is sent with a value less than ACCOUNT_UNLOCK_TIMEOUT = 600 seconds (default), 200 OK is returned instead of throwing an error.

Here is an example:

```
# curl -s -k -H "X-Auth-Token: $token" -X PATCH
https://192.168.31.1/redfish/v1/AccountService -d '{"AccountLockoutDuration":
300}'
200 Ok
{
  "@Message.ExtendedInfo": [
    {
      "@odata.type": "#Message.v1_1_1.Message",
      "Message": "The request completed successfully.",
      "MessageArgs": [],
      "MessageId": "Base.1.15.0.Success",
      "MessageSeverity": "OK",
      "Resolution": "None"
    }
  ]
}
# curl -k -H "X-Auth-Token: $token" -X GET https://192.168.31.1/redfish/v1/
AccountService
| grep AccountLockout
  "AccountLockoutDuration": 600,
  "AccountLockoutThreshold": 3,
```

600 seconds is the default waiting time before trying to log in again when the account is locked due to multiple login failures. This value cannot be changed to a value that is less than 600 seconds using PATCH request.

In the current implementation, the issue is happening because of the inability to handle values less than the compile time limit (for example, ACCOUNT_UNLOCK_TIMEOUT = 600 seconds). When a 200 OK response is returned, users think that they have updated the account timeout value, but the original behavior does not change.

Workaround

There is no workaround at this time.

- ▶ During a stress test, the HMC ERoT failed to get the SPDM measurements. – 4130662

During a stress test, where the firmware update and the SPDM attestation data are being completed simultaneously, the SPDM attestation API sometimes fails to get the SPDM Attestation data.

Workaround

Retrying the operation might result in a successful outcome. The fix will be included in the v1.3.0 GA release.

- ▶ Redfish Service Validator reports the following errors: – 4099508

- ▶ ERROR - 1 err.SessionCollection.SessionCollection errors in /redfish/v1/
- ▶ ERROR - 1 failMandatoryExist errors in /redfish/v1/
- ▶ ERROR - 1 failRedfishUriStrict errors in /redfish/v1/Managers/HGX_BMC_0
- ▶ ERROR - 1 failRedfishUriStrict errors in /redfish/v1/Managers/HGX_BMC_0/LogServices/Journal
- ▶ ERROR - 1 err.PCIeDevice.PCIeTypes errors in /redfish/v1/Systems/HGX_Baseboard_0/Processors/FPGA_0
- ▶ ERROR - 1 failProp errors in /redfish/v1/Systems/HGX_Baseboard_0/Processors/FPGA_0
- ▶ ERROR - 1 failRedfishUriStrict errors in /redfish/v1/Systems/HGX_Baseboard_0/LogServices/DebugTokenService
- ▶ ERROR - Validation has failed: 5 problems found
- ▶ The Id property does not match the last segment of the URI:
 - ▶ ERROR - 1 failRedfishUriStrict errors in /redfish/v1/Managers/HGX_BMC_0
 - ▶ ERROR - 1 failRedfishUriStrict errors in /redfish/v1/Managers/HGX_BMC_0/LogServices/Journal
 - ▶ ERROR - 1 failRedfishUriStrict errors in /redfish/v1/Systems/HGX_Baseboard_0/LogServices/DebugTokenService

Workaround There is no functional impact and no workaround at this moment.

- ▶ ERROR - PCIeType shows "empty string": Value Enum not found in ['Gen1', 'Gen2', 'Gen3', 'Gen4', 'Gen5']
- ▶ ERROR - 1 err.PCIeDevice.PCIeTypes errors in /redfish/v1/Systems/HGX_Baseboard_0/Processors/FPGA_0
- ▶ ERROR - 1 failProp errors in /redfish/v1/Systems/HGX_Baseboard_0/Processors/FPGA_0

Workaround Check the `CurrentSpeedGbps` property on the Port resource of the FPGA:

- ▶ /redfish/v1/Systems/HGX_Baseboard_0/Processors/FPGA_0/Ports/PCIeToHost_0
- ▶ /redfish/v1/Systems/HGX_Baseboard_0/Processors/FPGA_0/Ports/PCIeToHMC_0
- ▶ The PCIeSwitch ERoT Recovery event does not generate a corresponding event on the PCIeSwitch AP Chassis URI. – 4325634

When the PCIeSwitch ERoT Recovery event is generated, there is no corresponding event in the Conditions array of the /redfish/v1/Chassis/HGX_PCIeSwitch_0 URI.

Workaround

If the PCIeSwitch Health or HealthRollup shows `Critical`, customers can check the ERoT Chassis URI (`/redfish/v1/Chassis/HGX_ERoT_PCIESwitch_0`) for Conditions.

- ▶ 4333486 CUDA initialization failure on MIG and forward compatibility setups. – 4333486

CUDA will not work in Multi Instance GPU (MIG) mode on CUDA forward compatibility setups with display driver major version 470 or earlier.

CUDA will not initialize correctly when run in forward compatibility mode with display driver from major version 470 and earlier when MIG is enabled. The CUDA APIs will return `cudaErrorNoDevice/CUDA_ERROR_NO_DEVICE` errors when run on such a setup.

Workaround

Update the display driver to a major version, which is supported in CUDA forward compatibility and that is later than 470 (for example, version 510 or later).

- ▶ `AutoClearResolvedLogEnabled` is not defined in the Complex NVIDIA `NvidiaLogService.v1_0_0.NvidiaLogService`. 4313643

The following message is displayed in the Redfish service validator:

```
ERROR - 1 failAdditional.complex errors in /redfish/v1/Systems/HGX_Baseboard_0/LogServices/EventLog
```

This issue occurs on the `/redfish/v1/Systems/HGX_Baseboard_0/LogServices/EventLog` URI because the `Oem.Nvidia @odata.type` property has the `NvidiaLogService.v1_0_0.NvidiaLogService` value instead of the `NvidiaLogService.v1_1_0.NvidiaLogService` value.

Here is the sample output of the `/redfish/v1/Systems/HGX_Baseboard_0/LogServices/EventLog` URI:

```
{
  "@odata.id":
    "/redfish/v1/Systems/HGX_Baseboard_0/LogServices/EventLog",
  "@odata.type": "#LogService.v1_1_0.LogService",
  "Actions": {
    "#LogService.ClearLog": {
      "target":
        "/redfish/v1/Systems/HGX_Baseboard_0/LogServices/EventLog/Actions/LogService.ClearLog"
    }
  },
  "DateTime": "2020-08-23T23:51:08+00:00",
  "DateTimeLocalOffset": "+00:00",
  "Description": "System Event Log Service",
```

```

"Entries": {
  "@odata.id":
    "/redfish/v1/Systems/HGX_Baseboard_0/LogServices/EventLog/Entries"
},
"Id": "EventLog",
"Name": "Event Log Service",
"Oem": {
  "Nvidia": {
    "@odata.type":
      "#NvidiaLogService.v1_0_0.NvidiaLogService",
    "AutoClearResolvedLogEnabled": false,
    "LatestEntryID": "443",
    "LatestEntryTimeStamp":
      "2020-08-23T23:51:06+00:00"
  }
},
"OverWritePolicy": "WrapsWhenFull"
}

```

Workaround

There is no functional impact and no workaround at this time.

- ▶ No Secure boot failure event is generated. – 4299429

The HMC supports secure boot failure, SPI flash error, boot complete failure, heartbeat failure, ERoT fatal, ERoT recovery, and MCTP runtime error ERoT-related events at bootup and runtime. Issues with incoming firmware packages during the firmware update process are also reported by the HMC. A secure boot failure as the result of a key revocation will be detected when users attempt to flash a firmware update package with a revoked key. When this occurs, the following message is displayed in the firmware update task log:

```

{
  "@odata.type": "#MessageRegistry.v1_4_1.MessageRegistry",
  "Message": "The resource property 'HGX_FW_BMC_0' has detected errors of
    type
      'Security keys revoked'.",
  "MessageArgs": [
    "HGX_FW_BMC_0",
    "Security keys revoked"
  ]
}

```

```

],
  "MessageId": "ResourceEvent.1.0.ResourceErrorsDetected",
  "Resolution": "Verify the contents of the FW package",
  "Severity": "Critical"
}

```

However, when the image already on the AP flash is signed with a key that is revoked, no Secure boot failure event is generated. In this case the device Health/HealthRollup becomes Critical.

If the AP fails to boot because both images are signed with a revoked key, the EROt fatal event is raised on the associated EROt Chassis URI, and for GPU/NVSwitch a NACK From Device event on the AP Chassis URI is also seen. A secure boot failure as the result of a signature authentication failure is still reported during the firmware update and at boot or runtime. This scenario triggers the following Redfish event, which shows in the Conditions of the AP Chassis URI:

```

{
  "@odata.id":
    "/redfish/v1/Systems/HGX_Baseboard_0/LogServices/EventLog/Entries/2",
  "@odata.type": "#LogEntry.v1_9_0.LogEntry",
  "AdditionalDataURI":
    "/redfish/v1/Systems/HGX_Baseboard_0/LogServices/EventLog/Entries/2/attachment",
  "Created": "2034-09-11T23:22:25+00:00",
  "EntryType": "Event",
  "Id": "2",
  "Links": {
    "OriginOfCondition": {
      "@odata.id":
        "/redfish/v1/UpdateService/FirmwareInventory/HGX_FW_NVSwitch_0"
    }
  },
  "Message": "The resource property NVSwitch_0 Firmware has detected errors of type 'Secure Boot Failure'.",
  "MessageArgs": [
    "NVSwitch_0 Firmware",
    "Secure Boot Failure"
  ],

```

```
"MessageId": "ResourceEvent.1.0.ResourceErrorsDetected",
"Name": "System Event Log Entry",
  "Resolution": "ERoT will try to recover AP by a reset/reboot. If there is
    still a problem, collect ERoT logs and reflash AP firmware with recovery
    flow.",
"Resolved": false,
"Severity": "Critical"
},
```

- ▶ OpenCL is not supported in the confidential compute mode available with Hopper (<https://www.nvidia.com/en-in/data-center/solutions/confidential-computing/>) . Any attempt to run OpenCL apps in this mode will result in `clGetPlatformIDs` returning an error `CL_INVALID_DEVICE`.
- ▶ All Mellanox ports are shown by command `"nvidia-smi topo -m"` after dual-port NICs are bonded.
- ▶ When polling the H100 GPU via SMBPBI using GPU Performance Monitoring metrics, driver reloads or GPU resets can result in driver errors that manifest as PID (X62) errors on Linux. NVIDIA is investigating this issue.
- ▶ On NVIDIA H800, monitoring software such as DCGM or NVML might report lower double-precision (FP64) utilization metrics. This is expected as per the NVIDIA H800 product configuration. Refer to the NVIDIA H800 product brief for more details.
- ▶ The GPU driver build system might not pick the `Module.symvers` file, produced when building the `ofa_kernel` module from `MLNX_OFED`, from the right subdirectory. Because of that, `nvidia_peermem.ko` does not have the right kernel symbol versions for the APIs exported by the IB core driver, and therefore it does not load correctly. That happens when using `MLNX_OFED` 5.5 or newer on a Linux Arm64 or ppc64le platform.

To work around this issue, perform the following:

1. Verify that `nvidia_peermem.ko` does not load correctly.
2. Uninstall old `MLNX_OFED` if one was installed.
3. Manually remove `/usr/src/ofa_kernel/default` if one exists.
4. Install `MLNX_OFED` 5.5 or newer.
5. Manually create a soft link:

```
/usr/src/ofa_kernel/default -> /usr/src/ofa_kernel/$(uname -m)/$(uname -r)
```

6. Reinstall the GPU driver.

- ▶ If you encounter an error on RHEL7 when installing with `cuda-drivers-fabricmanager` packages, use the following alternate instructions. For example:

If you are upgrading from a different branch, for example to driver 515.65.01:

```
new_version=515.65.01
sudo yum swap nvidia-driver-latest-dkms nvidia-driver-latest-dkms-${new_version}
```

```
sudo yum install nvidia-fabric-manager- $\${new\_version}$ 
```

- ▶ When installing a driver on SLES15 or openSUSE15 that previously had an R515 driver installed, users need to run the following command afterwards to finalize the installation:

```
sudo zypper install --force nvidia-gfxG05-kmp-default
```

Without doing this, users may see the kernel objects as missing.

- ▶ `nvidia-release-upgrade` may report that not all updates have been installed and exit.

When running the

```
nvidia-release-upgrade
```

command on DGX systems running DGX OS 4.99.x, it may exit and tell users: "Please install all available updates for your release before upgrading" even though all upgrades have been installed.

Users who see this can run the following command:

```
sudo apt install -y nvidia-fabricmanager-450/bionic-updates --allow-downgrades
```

After running this, proceed with the regular upgrade steps:

```
sudo apt update
sudo apt full-upgrade -y
sudo apt install -y nvidia-release-upgrade
sudo nvidia-release-upgrade
```

- ▶ By default, Fabric Manager runs as a `systemd` service. If using `DAEMONIZE=0` in the Fabric Manager configuration file, then the following steps may be required.

1. Disable FM service from auto starting.

```
systemctl disable nvidia-fabricmanager
```

2. Once the system is booted, manually start FM process.

```
/usr/bin/nv-fabricmanager -c /usr/share/nvidia/nvswitch/fabricmanager.cfg
```

Note, since the process is not a daemon, the SSH/Shell prompt will not be returned (use another SSH shell for other activities or run FM as a background task).

- ▶ Important correctness fix for H100 GPU instructions used by cuBLAS, other CUDA libraries, and user CUDA code

An issue was discovered recently with H100 GPUs (H100 PCIe and HGX H100) where certain operations put the GPU in an invalid state that allowed some GPU instructions to operate at unsupported frequency that can result in incorrect computation results and faster than expected performance. The affected GPU instructions are used by cuBLAS, other CUDA libraries, and can also be used for user CUDA code.

The operations that allow the GPU to enter an invalid state are the following:

- ▶ Enabling MIG
- ▶ Deinitialize and reinitialize the GPU (for example, turn off persistence mode and turn it back on or reload the nvidia.ko driver)
- ▶ Any Compute Engine error (for example, MMU fault, Out of Range warp error, and so on)

Once the GPU enters the invalid state, the performance for some GPU instructions is increased by 7-10%, but the computation results may be incorrect.

The current release fixes this issue, and it is no longer possible to enter the invalid GPU state. This issue has been present in all drivers since the H100 launch, and we recommend that you upgrade to the current release as soon as possible. If upgrading is not immediately possible, a GPU reset can restore the GPU back to the correct operational state, except for when MIG is being used. For MIG, the new driver is required, and there is no workaround available.

- ▶ Uninstalling the driver fails, and the system reboots automatically.

On Windows 2019 and 2022 servers, uninstalling the driver causes the system to restart automatically before the uninstallation is completed. The issue also occurs when you upgrade the driver from an older version to a new version, even after selecting the Perform Clean Installation option in the installer UI.



Note: This issue does not occur in Linux.

Workaround

We strongly recommend that you always install, uninstall, and upgrade drivers from Safe mode.

- ▶ In Shared Switch virtualization mode, the guest VM GPU driver load and unload stress test fails after certain iteration

In the Shared Switch virtualization mode, the stress test to load and unload the GPU driver on Guest VM in every 30 second interval runs into issues approximately after three hours of the test.

Workaround

Do not run the stress reload driver cycle at this time.

- ▶ A few Async SMBPBI commands do not function as intended when the driver is unloaded.

When the driver is unloaded, the following Async SMBPBI commands do not operate as specified:

- ▶ Arg1 0x00: Reads total GPU power limit control data.
- ▶ Arg1 0x01: Sets the total GPU power limit.
- ▶ Arg1 0x02: Reads the total GPU power limit policy information.

Due to this issue, some properties of the following Redfish URIs are impacted:

- ▶ **PowerLimitWatts.SetPoint:**

```
/redfish/v1/Systems/HGX_Baseboard_0/Processors/GPU_SXM_[1-8]/
EnvironmentMetrics
```

- ▶ **SpeedLimitMHz, SpeedLocked:**

```
/redfish/v1/Systems/HGX_Baseboard_0/Processors/GPU_SXM_[1-8]
```

The Patch operation of the following URIs are impacted:

- ▶ **PowerLimitWatts.SetPoint:**

```
/redfish/v1/Systems/HGX_Baseboard_0/Processors/GPU_SXM_[1-8]/
EnvironmentMetrics
```

- ▶ **Oem.Nvidia.PowerMode "MaxP" or "MaxQ":**

```
/redfish/v1/Chassis/HGX_Chassis_0/EnvironmentMetrics
```

- ▶ **SpeedLimitMHz, SpeedLocked:**

```
/redfish/v1/Systems/HGX_Baseboard_0/Processors/GPU_SXM_[1-8]
```

Workaround

Load the driver for these URIs to work properly.

- ▶ Fabric Manager state is not reported accurately on NVSwitch OOB query

The NVSwitch SMPBI query that reports Fabric Manager state (Manager State) is not reporting the actual FM state.

- ▶ Instructions to reset all GPUs Using the `nvidia-smi -r` Command

When resetting all GPUs using the `nvidia-smi` command with the `-r` option instead of a resetting specific GPU using the `-i <gpu_index>` option, all the NVSwitches will also be reset. This process wipes out the NVSwitch routing entries, and subsequent CUDA application launches will fail. The Fabric Manager service will also show interaction errors with the NVSwitch device via the switch driver.

Workaround

1. Stop the Fabric Manager service.
2. To reset all GPUs, run `nvidia-smi -r`.
3. After the reset is finished, start the Fabric Manager service.

GPU Performance Counters

The use of developer tools from NVIDIA that access various performance counters requires administrator privileges. See this [note](#) for more details. For example, reading NVLink utilization metrics from `nvidia-smi` (`nvidia-smi nvlink -g 0`) would require administrator privileges.

NoScanout Mode

NoScanout mode is no longer supported on NVIDIA Data Center GPU products. If NoScanout mode was previously used, then the following line in the “screen” section of `/etc/X11/xorg.conf` should be removed to ensure that X server starts on data center products:

```
Option          "UseDisplayDevice" "None"
```

NVIDIA Data Center GPU products now support one display of up to 2560x1600 resolution.

Unified Memory Support

CUDA and unified memory is not supported when used with Linux power management states S3/S4.

IMPV FRU for Volta GPUs

The driver does not support the IPMI FRU multi-record information structure for NVLink. See the Design Guide for Tesla P100 and Tesla V100-SXM2 for more information.

OpenCL 3.0 Known Issues

Device side enqueue

- ▶ Device-Side-Enqueue related queries may return 0 values, although corresponding built-ins can be safely used by kernel. This is in accordance with conformance requirements described at https://www.khronos.org/registry/OpenCL/specs/3.0-unified/html/OpenCL_API.html#opencl-3.0-backwardscompatibility
- ▶ Shared virtual memory - the current implementation of shared virtual memory is limited to 64-bit platforms only.

Confidential Compute

- ▶ Confidential Compute applications will not work on this release. Please continue to use 535.104.05 for this use case.

Chapter 2. Virtualization

To make use of GPU passthrough with virtual machines running Windows and Linux, the hardware platform must support the following features:

- ▶ A CPU with hardware-assisted instruction set virtualization: Intel VT-x or AMD-V.
- ▶ Platform support for I/O DMA remapping.
- ▶ On Intel platforms, the DMA remapper technology is called Intel VT-d.
- ▶ On AMD platforms, it is called AMD IOMMU.

Support for these features varies by processor family, product, and system, and should be verified at the manufacturer's website.

The following hypervisors are supported for virtualization:

Hypervisor	Notes
Citrix XenServer	Version 6.0 and later
VMware vSphere (ESX / ESXi)	Version 5.1 and later.
Red Hat KVM	Red Hat Enterprise Linux 7 with KVM
Microsoft Hyper-V	Windows Server 2019 Hyper-V Generation 2

Data Center products now support one display of up to 2560x1600 resolution.

The following GPUs are supported for device passthrough for virtualization:

GPU Family	Boards Supported
NVIDIA Ada Lovelace	NVIDIA L40, L4
NVIDIA Grace Hopper	NVIDIA GH200
NVIDIA Hopper	NVIDIA H100, NVIDIA H800
NVIDIA Ampere GPU Architecture	NVIDIA A800, A100, A40, A30, A16, A10, A10G, A2, AX800
NVIDIA Turing	NVIDIA T4, NVIDIA T4G
NVIDIA Volta	NVIDIA V100
NVIDIA Pascal	Quadro: P2000, P4000, P5000, P6000, GP100

GPU Family	Boards Supported
NVIDIA Maxwell	Tesla: P 100, P40, P4 Quadro: K2200, M2000, M4000, M5000, M6000, M6000 24GB Tesla: M60, M40, M6, M4

Chapter 3. Hardware and Software Support

Support for these features varies by processor family, product, and system, and should be verified at the manufacturer's website.

Supported Operating Systems for NVIDIA Data Center GPUs

The Release 535 driver is supported on the following operating systems:

- ▶ Windows x86_64 operating systems:
 - ▶ Microsoft Windows® Server 2022
 - ▶ Microsoft Windows® Server 2019
 - ▶ Note: R525TeslaRD was the last TRD to support Server 2016.
 - ▶ Microsoft Windows® 11 21H2 - SV1
 - ▶ Microsoft Windows® 11 22H2 - SV2
 - ▶ Microsoft Windows® 10
- ▶ The following table summarizes the supported Linux 64-bit distributions. For a complete list of distributions, kernel versions supported, see the [CUDA Linux System Requirements](#) documentation.

Distribution	x86_64	POWER	Arm64 Server
Debian 11.x (where x <= 7)	Yes	No	No
Debian 10. x (where x <= 13)	Yes	No	No
OpenSUSE Leap 15.x (where y <= 5)	Yes	No	No
Fedora 37	Yes	No	No
Red Hat Enterprise Linux 9.y (where y <= 2)	Yes	No	Yes

Distribution	x86_64	POWER	Arm64 Server
Rocky Linux 9.y (where y <= 2)	Yes	No	No
Red Hat Enterprise Linux 8.y (where y <= 8)	Yes	Yes	Yes
Rocky Linux 8.y (where y <= 8)	Yes	No	No
Red Hat Enterprise Linux / CentOS 7.y (where y <= 9)	Yes	No	No
SUSE Linux Enterprise Server 15.y (where y <= 5)	Yes	No	Yes
Ubuntu 22.04.z LTS (where z <= 3)	Yes	No	Yes
Ubuntu 20.04.z LTS (where z <= 6)	Yes	No	Yes
KylinOS V10 SP2	Yes	No	No
CBL-Mariner 2.0*	Yes	No	No

* CBL-Mariner will be supported by TRD via runfile. CUDA Toolkit will not support this OS as this is a deployment OS.

Supported Operating Systems and CPU Configurations for NVIDIA HGX GH200

- ▶ Grace Hopper Linux distributions:
 - ▶ Red Hat Enterprise Linux 9.3
 - ▶ SUSE Linux Enterprise Server 15 SP5 QU1
 - ▶ Ubuntu 22.04 with NVIDIA HWE kernel

RHEL and SLES feature parity with NVIDIA HWE Kernels. The latest RHEL 9 and SLES 15 SP5 kernels support bare metal.

Supported Operating Systems and CPU Configurations for NVIDIA HGX H100/H800

The Release 535 driver is validated with NVIDIA HGX H100 on the following operating systems and CPU configurations:

- ▶ Linux 64-bit distributions:

- ▶ Red Hat Enterprise Linux 8.8 (in 4/8/16-GPU configurations)
- ▶ Red Hat Enterprise Linux 9.2 (in 4/8/16-GPU configurations)
- ▶ SUSE Linux Enterprise Server 15.5 (in 4/8/16-GPU configurations)
- ▶ Ubuntu 22.04.2 LTS (in 4/8/16-GPU configurations)
- ▶ Windows 64-bit distributions:
 - ▶ Windows Server 2022
 - ▶ Windows Server 2019 (in 1/2/4/8-GPU configurations; 16-GPU configurations are currently not supported)

Windows is supported only in shared NVSwitch virtualization configurations.

Supported Operating Systems and CPU Configurations for NVIDIA HGX A100/A800

The Release 535 driver is validated with NVIDIA HGX A100 on the following operating systems and CPU configurations:

- ▶ Linux 64-bit distributions:
 - ▶ Debian 11.7
 - ▶ Debian 10.13
 - ▶ Red Hat Enterprise Linux 8.8 (in 4/8/16-GPU configurations)
 - ▶ Red Hat Enterprise Linux 7.9 (in 4/8/16-GPU configurations)
 - ▶ Rocky Linux 8.8 (in 4/8/16-GPU configurations)
 - ▶ Red Hat Enterprise Linux 9.2 (in 4/8/16-GPU configurations)
 - ▶ CentOS Linux 7.9 (in 4/8/16-GPU configurations)
 - ▶ Ubuntu 22.04.3 LTS (in 4/8/16-GPU configurations)
 - ▶ Ubuntu 20.04.6 LTS (in 4/8/16-GPU configurations)
 - ▶ SUSE SLES 15.5 (in 4/8/16-GPU configurations)
 - ▶ KylinOS V10 SP2
- ▶ Windows 64-bit distributions:
 - ▶ Windows Server 2022
 - ▶ Windows Server 2019 (in 1/2/4/8-GPU configurations; 16-GPU configurations are currently not supported)

Windows is supported only in shared NVSwitch virtualization configurations.
- ▶ CPU Configurations:
 - ▶ AMD Rome in PCIe Gen4 mode
 - ▶ Intel Skylake/Cascade Lake (4-socket) in PCIe Gen3 mode

Supported Virtualization Configurations

The Release 535 driver is validated with NVIDIA HGX A100, HGX A800, H100, and H800 on the following configurations:

- ▶ Passthrough (full visibility of GPUs and NVSwitches to guest VMs):
 - ▶ 8-GPU configurations with Ubuntu 20.04.6 and 22.4.2
- ▶ Shared NVSwitch (guest VMs only have visibility of GPUs and full NVLink bandwidth between GPUs in the same guest VM):
 - ▶ 1/2/4/8/16-GPU configurations with Ubuntu 20.04.6 LTS

API Support

This release supports the following APIs:

- ▶ NVIDIA® CUDA® 12.2 for NVIDIA® Maxwell™, Pascal™, Volta™, Turing™, Hopper™, NVIDIA Ampere architecture, and NVIDIA Ada Lovelace GPU architecture GPUs
- ▶ OpenGL® 4.6
- ▶ Vulkan® 1.3
- ▶ DirectX 11
- ▶ DirectX 12 (Windows 10)
- ▶ Open Computing Language (OpenCL™ software) 3.0

Note that for using graphics APIs on Windows (such as OpenGL, Vulkan, DirectX 11, and DirectX 12) or any WDDM 2.0+ based functionality on Data Center GPUs, vGPU is required. See the [vGPU documentation](#) for more information.

Supported NVIDIA Data Center GPUs

The NVIDIA Data Center GPU driver package is designed for systems that have one or more Data Center GPU products installed. This release of the driver supports CUDA C/C++ applications and libraries that rely on the CUDA C Runtime and/or CUDA Driver API.

Attention: Release 470 was the last driver branch to support Data Center GPUs based on the NVIDIA Kepler architecture. This includes discontinued support for the following compute capabilities:

- ▶ sm_30 (NVIDIA Kepler)
- ▶ sm_32 (NVIDIA Kepler)
- ▶ sm_35 (NVIDIA Kepler)
- ▶ sm_37 (NVIDIA Kepler)

For more information on GPU products and compute capability, see <https://developer.nvidia.com/cuda-gpus>.

NVIDIA Server Platforms	
Product	Architecture
NVIDIA HGX H100	H100 and NVSwitch
NVIDIA HGX H800	H800 and NVSwitch
NVIDIA HGX A800	A800 and NVSwitch
NVIDIA HGX A100	A100 and NVSwitch
NVIDIA HGX-2	V100 and NVSwitch

Data Center L-Series Products	
Product	GPU Architecture
NVIDIA L40	NVIDIA Ada Lovelace
NVIDIA L40S	NVIDIA Ada Lovelace
NVIDIA L4	NVIDIA Ada Lovelace

Data Center H-Series Products	
Product	GPU Architecture
NVIDIA H100 PCIe	NVIDIA Hopper
NVIDIA H100 NVL	NVIDIA Hopper
NVIDIA H800 PCIe	NVIDIA Hopper
NVIDIA H800 NVL	NVIDIA Hopper

RTX-Series / T-Series Products	
Product	GPU Architecture
NVIDIA RTX 6000 Ada Generation	NVIDIA Ada Lovelace
NVIDIA RTX 4000 SFF Ada Generation	NVIDIA Ada Lovelace
NVIDIA RTX A6000	NVIDIA Ampere architecture
NVIDIA RTX A5000	NVIDIA Ampere architecture
NVIDIA RTX A4000	NVIDIA Ampere architecture
Quadro RTX 8000	NVIDIA Turing
Quadro RTX 6000	NVIDIA Turing
Quadro RTX 4000	NVIDIA Turing
NVIDIA T1000	NVIDIA Turing
NVIDIA T600	NVIDIA Turing
NVIDIA T400	NVIDIA Turing

Data Center A-Series Products	
Product	GPU Architecture
NVIDIA A2	NVIDIA Ampere architecture
NVIDIA A800, AX800	NVIDIA Ampere architecture
NVIDIA A100X	NVIDIA Ampere architecture
NVIDIA A100	NVIDIA Ampere architecture
NVIDIA A100 80 GB PCIe	
NVIDIA A40	NVIDIA Ampere architecture
NVIDIA A30, A30X	NVIDIA Ampere architecture
NVIDIA A16	NVIDIA Ampere architecture
NVIDIA A10, A10M, A10G	NVIDIA Ampere architecture

Data Center T-Series Products	
Product	GPU Architecture
NVIDIA T4, T4G	NVIDIA Turing

Data Center V-Series Products	
Product	GPU Architecture
NVIDIA V100	Volta

Data Center P-Series Products	
Product	GPU Architecture
NVIDIA Tesla P100	NVIDIA Pascal
NVIDIA Tesla P40	NVIDIA Pascal
NVIDIA Tesla P4	NVIDIA Pascal

Data Center M-Class Products	
Product	GPU Architecture
NVIDIA Tesla M60	Maxwell
NVIDIA Tesla M40 24 GB	Maxwell
NVIDIA Tesla M40	Maxwell
NVIDIA Tesla M6	Maxwell
NVIDIA Tesla M4	Maxwell

Notice

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

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

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

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

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

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

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

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

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

Trademarks

NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the United States and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2023-2024 NVIDIA Corporation and affiliates. All rights reserved.