



# NVIDIA Magnum IO GPUDirect Storage

Release Notes

# Table of Contents

Chapter 1. Introduction.....	1
Chapter 2. New Features.....	2
Chapter 3. MLNX_OFED and Filesystem Requirements.....	3
Chapter 4. Support Matrix.....	4
Chapter 5. Included Packages.....	5
Chapter 6. Minor Updates and Bug Fixes.....	6
Chapter 7. Known Issues.....	7
Chapter 8. Known Limitations.....	8

---

# Chapter 1. Introduction

Release information for NVIDIA® GPUDirect® Storage (GDS) for developers and users.

NVIDIA® Magnum IO GPUDirect® Storage (GDS) is one of the members of the GPUDirect family of technologies. GDS enables a direct data path for direct memory access (DMA) transfers between GPU memory and storage. This direct path increases IO bandwidth, decreases IO latency and reduces the utilization load on the host CPU.

GDS is generally available on third party storage solutions from DDN EXAScaler®, Dell Isilon, IBM Spectrum Scale, VAST NFS, WekaFS™, NVmesh Excelero, ScaleFlux, and Micron. GDS documents and online resources provide additional context for the optimal use of and understanding of GPUDirect Storage.

Refer to the following guides for more information about GDS:

- ▶ [GPUDirect Storage Design Guide](#)
- ▶ [GPUDirect Storage Overview Guide](#)
- ▶ [cuFile API Reference Guide](#)
- ▶ [GPUDirect Storage Best Practices Guide](#)
- ▶ [GPUDirect Storage Installation and Troubleshooting Guide](#)
- ▶ [GPUDirect Storage O\\_DIRECT Requirements Guide](#)

To learn more about GDS, refer to the following posts:

- ▶ [GPUDirect Storage: A Direct Path Between Storage and GPU Memory](#)
- ▶ The [Magnum IO](#) blog series.

---

## Chapter 2. New Features

The following features have been added in v1.1:

- ▶ The XFS file system has been added to the list of supported file systems at a beta support level.
- ▶ Improved support for unregistered buffers.
- ▶ Added options `start_offset` and `io_size` to `gdsio` config file per job options.
- ▶ Improved performance of 4K and 8K IO sizes for local file systems.
- ▶ Added user-configurable priority for internal `cuFile` CUDA streams.

---

# Chapter 3. MLNX\_OFED and Filesystem Requirements

The following are the MLNX\_OFED and filesystem requirements for GDS:

- ▶ MLNX\_OFED 5.3 and later, which supports NVMe NVMeoF, NFSoRDMA (VAST) on Linux kernel 4.15.x and 5.4.x. MLNX\_OFED 5.4 is preferred.

MLNX\_OFED must be installed **before** installing GDS. Refer to [Installing GPUDirect Storage](#) for more information about installing MLNX\_OFED.

- ▶ The following distributed filesystems are supported:
  - ▶ DDN EXAScaler 5.2
  - ▶ IBM Spectrum Scale
  - ▶ VAST 3.4
  - ▶ WekaFS 3.8.0
  - ▶ Dell Powerscale v9.2.0.0
- ▶ Block/other file systems supported:
  - ▶ ExceleroNVMesh 2.2
  - ▶ Pavilion Data 2.0 or later
  - ▶ ScaleFlux CSD 3.2.7.0

---

# Chapter 4. Support Matrix

GDS has been tested on following NVIDIA GPUs: T10x, T4, A10, Quadro P6000, A100, and V100.

	<b>GDS mode supported (includes fallback compat mode)</b>	<b>Only compatibility mode supported</b>
<b>x86-64</b>	RHEL 8.3, 8.4. Ubuntu 18.04,20.04	Debian 10, OpenSUSE 15.2, RHEL 7.9, CentOS 8.3, SLES 15.2

---

# Chapter 5. Included Packages

The GDS package contains the following Debian packages:

- ▶ `gds-tools-11-5_*.deb`
- ▶ `libcufile-11-5_*.deb`
- ▶ `libcufile-dev-11-5_*.deb`
- ▶ `nvidia-fs_2.8*.deb`
- ▶ `nvidia-fs-dkms_2.8*.deb`
- ▶ `nvidia-gds-11-5_*.deb`
- ▶ `nvidia-gds_11.5.*.deb`



**Note:** Each component has a README file. For example, for `gds-tools`, the README file is in the `/usr/local/CUDA-11.5/gds/tools/` directory.

---

## Chapter 6. Minor Updates and Bug Fixes

The following minor updates and bug fixes were made in version 1.1:

- ▶ Fixed crash in libcufile when BAR1 space is unavailable.
- ▶ Fixed error stats reporting in nvidia-fs.
- ▶ Fixed read failures in lustre for files less than 4K when GPU buffer is not 64K aligned.



---

# Chapter 7. Known Issues

This section provides information about the known issues in this release of GDS.

- ▶ If you see an error in cuFile “`ERROR cufile:671 assertion: fd >= 0`”, it can be ignored. This issue will be fixed in an upcoming release.
- ▶ If the volume/file system where `cufile.log` is residing runs out of space, then a GPU application using GDS is likely to fail and exit. This issue will be fixed in an upcoming release.
- ▶ For DDN EXAScaler filesystem:
  - ▶ With stripe count > 1, `cuFileRead` and `cuFileWrite` do not work with poll mode enabled for versions older than 2.12.5\_ddn10.
  - ▶ With 2.12.5\_ddn10, any reads beyond EOF causes a `BUG_ON` inside `nvidia-fs`.
- ▶ RHEL8.3 does not have default udev rules for detecting RAID members, which disables GDS on RAID volumes. Refer to the section [Adding udev Rules for RAID Volumes](#) in the GPUDirect Storage Installation and Troubleshooting Guide.
- ▶ On DGX OS:
  - ▶ For log collection, use `gds_log_collection.py` described in [Sending Relevant Data to Customer Support](#).
  - ▶ For package installation for releases earlier than CUDA 11.4, use the online instructions in [Preparing the OS](#) instead of the README included as part of the GDS package.

---

# Chapter 8. Known Limitations

This section provides information about the known limitations in this release of GDS.

- ▶ Data Center and Quadro (desktop) cards with compute capability > 6 listed here - <https://developer.nvidia.com/cuda-gpus#compute> are supported in GDS mode. All other cards are supported only in compatibility mode.
- ▶ With any NVIDIA software components installed, downgrading from RHEL 8.4 to RHEL 8.3 is not supported.
- ▶ For DDN EXAScaler, checksum is disabled in the read/write IO path.
- ▶ For WekaFS, checksum is disabled in the read/write IO path.
- ▶ cuFile APIs are not supported with applications using the `fork()` system call.
- ▶ GDS Compatibility mode works on GDS qualified file systems: EXAScaler, XFS, ext4, IBM Spectrum Scale, VAST, and WekaFS.
- ▶ GDS with IOMMU enabled or ACS enabled are not guaranteed to work functionally or in a performant way with all non-DGX based platforms.
- ▶ Refer to the following documentation for IBM Spectrum Scale Limitations with GDS: <http://www.ibm.com/support/pages/node/6444075>

## 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.



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.

## **HDMI**

HDMI, the HDMI logo, and High-Definition Multimedia Interface are trademarks or registered trademarks of HDMI Licensing LLC.

## **OpenCL**

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.



## Trademarks

NVIDIA, the NVIDIA logo, DGX, DGX-1, DGX-2, DGX-A100, Tesla, and Quadro 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

© 2021 NVIDIA Corporation and affiliates. All rights reserved.

