



NVIDIA Magnum IO GPUDirect Storage

Getting Started

Table of Contents

Chapter 1. Introduction.....	1
Chapter 2. If you are a system administrator or a performance engineer.....	2
Chapter 3. If you are a developer.....	3
Chapter 4. If you are OEM, ODM, CSP.....	5
Chapter 5. Troubleshooting GDS issues.....	6

Chapter 1. Introduction

NVIDIA® GPUDirect® Storage (GDS) is the newest addition to the GPUDirect family. GDS enables a direct data path for direct memory access (DMA) transfers between GPU memory and storage, which avoids a bounce buffer through the CPU. This direct path increases system bandwidth and decreases the latency and utilization load on the CPU.

The purpose of this guide is to help the user evaluate and test GDS functionality and performance by using sample applications. These applications can be run after you set up and install GDS and before you run the custom applications that have been modified to take advantage of GDS.

Refer to the following guides for more information about GDS:

- ▶ [GPUDirect Storage Design Guide](#)
- ▶ [GPUDirect Storage Overview Guide](#)
- ▶ [cuFile API Reference Guide](#)
- ▶ [GPUDirect Storage Release Notes](#)
- ▶ [GPUDirect Storage Best Practices Guide](#)
- ▶ [GPUDirect Storage Installation and Troubleshooting Guide](#)
- ▶ [GPUDirect Storage Benchmarking and Configuration 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](#) series.

Chapter 2. If you are a system administrator or a performance engineer

GDS is supported on systems which have a NIC/NVMe and GPU.

- ▶ The [GPUDirect Storage Design Guide](#) goes into the details of system architecture that will give the best performance.
- ▶ The [Support matrix](#) lists supported file systems.
- ▶ [File system settings](#) and performance tuning are covered in the *Installation and Troubleshooting Guide*.
- ▶ Application setup and performance benchmarkin are covered in the [GPUDirect Storage Benchmarking and Configuration Guide Guide](#).

Chapter 3. If you are a developer

Prior to starting application development with cuFile, it is highly recommended to go through the following information:

- ▶ [cuFile API Reference Guide](#)
- ▶ [GPUDirect Storage Best Practices Guide](#)
- ▶ Use of compatibility for development.

In order to develop applications, you do not need to complete all the installation steps to enable GDS. You can use existing CUDA toolkit installation from version 11.4 onward on any system that contains a data center class GPU. To enable compatibility mode, refer to [GDS and Compatibility Mode](#).

- ▶ How to build a GDS-enabled container:
 - ▶ In order to use GDS in a container, `nvidia-fs` must be installed on the host system. For example on Ubuntu:

```
$ sudo apt-get install nvidia-gds-12-1
```

- ▶ CUDA devel containers from [NGC](#) have all the necessary userspace libraries for GDS application development. The CUDA 12.1 devel container for Ubuntu 20.04 can be pulled with the following command:

```
$ sudo docker pull nvcr.io/nvidia/cuda:12.1.1-devel-ubuntu20.04
```

- ▶ The container should be started with the following docker run flags:
 - ▶ `--ipc host`
 - ▶ `--volume /run/udev:/run/udev:ro`
 - ▶ `--volume /sys/kernel/config:/sys/kernel/config:ro`
 - ▶ `--volume /dev:/dev:ro`
 - ▶ `--volume /lib/modules/$(uname -r)/:/lib/modules/$(uname -r):ro`
 - ▶ `--privileged`
- ▶ Applications and their IO characteristics.

In order to get best performance, two important parameters to consider are:

- ▶ Size of IO. For IO operations smaller than 64KB, cuFile synchronous APIs do not provide any performance benefits. It is strongly recommended to use Batch APIs instead.
- ▶ Whether IO is aligned or not.

- ▶ [IO patterns and sample code for each](#)
- ▶ [Samples and application benchmarks](#)
- ▶ [Python/C++ interface](#)

Chapter 4. If you are OEM, ODM, CSP

The [GPUDirect Storage Design Guide](#) describes how to design servers to take advantage of GPUDirect Storage.

Chapter 5. Troubleshooting GDS issues

The [GPUDirect Storage Installation and Troubleshooting Guide](#) has separate sections for basic and advanced troubleshooting.

The [Cheat Sheet for Diagnosing Problems](#) covers commonly seen problems and how to solve them.

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.

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

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

