

SupremeRAID™ User Guide for Linux

October, 2023



Copyright © 2021–2023 Graid Technology Inc. All Rights Reserved.

SupremeRAID™ is a registered trademark of Graid Technology Inc. All other trademarks and registered trademarks are the property of their respective owners.

Graid Technology reserves the right to make changes without further notice to any products or content herein to improve reliability, function, or design. Graid Technology makes no warranty as to the accuracy or completeness of the content or information provided herein, which are provided on an “as is” basis.

No license to Graid Technology’s or any third party’s intellectual property rights are conveyed hereunder.

Publication: August 2, 2023

CHANGE HISTORY

| Revision No. | Date | Prepared by /Modified by | Description |
|--------------|-------------------|--------------------------|---|
| 1.0.0 | November 20, 2022 | Liam Chen | Initial Draft |
| 1.0.1 | September 8, 2023 | Mia Lin | Update User Guide for 1.5.0 Linux Driver |
| 1.0.2 | October 23, 2023 | Mia Lin | Add Silent Installation and Background Task |

TABLE OF CONTENTS

| | |
|--|----|
| CHANGE HISTORY | 3 |
| INTRODUCTION..... | 10 |
| Software Module Overview | 10 |
| SupremeRAID™ Specification..... | 11 |
| RAID Components..... | 12 |
| Physical Drive (PD)..... | 13 |
| Drive Group (DG) | 13 |
| Virtual Drive (VD)..... | 13 |
| Features Overview | 14 |
| Ensuring Data Integrity with Consistency Checks..... | 14 |
| SupremeRAID's Dual-Controller Architecture for Auto-Failover and High-Availability | 15 |
| Setting Up the NVMe-oF Initiator Server and Managing Your RAID Components | 16 |
| Sharing the SupremeRAID™ Volume as a NVMe-oF Target Server | 17 |
| SPDK BDEV Feature of SupremeRAID™ | 18 |
| INSTALLATION | 19 |
| Prerequisites..... | 19 |
| Supported Operating Systems..... | 20 |
| Tested NVMe Devices | 22 |
| Installing the Hardware | 24 |
| ESD Warning | 24 |
| Installation Procedure..... | 24 |
| Installing the Software Driver..... | 26 |
| Using the Pre-installer and Installer..... | 26 |
| Using Installer for Silent Installation | 29 |
| Manual Installation..... | 31 |

Dependency Table for Manual Installation 31

Executing the Installer and Completing the Installation 43

USING THE SUPREMERAIID™ DRIVER 45

 Activating the SupremeRAID™ Driver and Managing the License(s) 45

 Example..... 46

 Creating a RAID-5 Virtual Drive with Five NVMe SSDs..... 47

 Output Example..... 47

 Creating a Physical Drive from the Remote NVMe-oF Targets..... 48

 Output Example..... 48

 Replace the Nearly Worn-out or Broken SSD. 49

 Output Example..... 49

 Exporting the Virtual Drive as an NVMe-oF Target Drive Using RDMA to the Initiator .. 50

 Output Example..... 50

 Setting Up the Dual-Controller to Enable HA and Auto-Failover..... 51

 Output Example..... 52

 Upgrading the Software..... 53

 Replacing a SupremeRAID™ Card..... 56

COMMANDS AND SHORTCUTS 58

 Syntax..... 58

 Command and Subcommand Quick Reference 59

 General..... 59

 Resources..... 59

 Features..... 61

 Managing Licenses..... 62

 Applying the License..... 62

 Checking License Information 63

 Checking the SupremeRAID™ Driver Version 65

 Output Example..... 65

 Viewing Host Drive Information..... 65

- Listing NVMe Drives..... 65
- Listing SAS/SATA Drives 67
- Managing Physical Drives..... 68
 - Creating a Physical Drive..... 68
 - Listing the Physical Drives..... 69
 - Deleting a Physical Drive..... 72
 - Describing a Physical Drive..... 72
 - Locating a Physical Drive..... 73
 - Marking a Physical Drive Online or Offline..... 73
 - Assigning a Hot Spare Drive 74
 - Replacing a Nearly Worn-Out or Broken SSD..... 74
- Managing Drive Groups..... 75
 - Creating Drive Groups..... 75
 - Deleting Drive Groups..... 79
 - Displaying Drive Group Information..... 80
 - Selecting the Controller for a Drive Group..... 83
 - Assigning a Controller to a Drive Group 83
 - Managing Background Task Speed 84
 - Locating the Physical Drives in the Drive Group..... 84
 - Degradation and Recovery 84
 - Rescue Mode..... 84
- Managing Virtual Drives..... 85
 - Creating a Virtual Drive 85
 - Listing Virtual Drives..... 86
 - Deleting Virtual Drives..... 88
 - Displaying Virtual Drive Information..... 89
 - Setting Up a Stripe Cache 89
- Managing Controllers 91
 - Activating a Controller..... 91

- Deactivating a Controller 91
- Listing Controllers..... 92
- Deleting a Controller 93
- Replacing a Controller License Key 94
- Importing and Controlling MD Bootable NVMe RAID1 95
 - Importing an MD Bootable NVMe RAID..... 95
 - Replacing an MD Bootable NVMe RAID1 96
 - Dismissing an Imported MD Bootable NVMe RAID1 98
- Adjusting or Updating Configuration Settings for the SupremeRAID™ Add-on 98
 - Editing Configuration Settings 98
 - Describing Configuration Settings..... 99
 - Deleting Configuration Settings 100
 - Restoring SupremeRAID™ Configuration Settings..... 101
- Managing Events 102
 - Listing Events 102
 - Deleting Events 103
- Managing Remote NVMe-oF Targets..... 103
 - Connecting to a Remote NVMe-oF Target 103
 - Listing Connected Remote NVMe-oF Targets..... 104
 - Disconnecting from Remote NVMe-oF Targets..... 105
- Exporting NVMe-oF Target Management..... 106
 - Creating the NVMe-oF Target Port Service 106
 - Exporting NVMe-oF Targets 107
 - Listing Created NVMe-oF Targets..... 108
 - Deleting the NVMe-oF Target Port Service Unexporting NVMe-oF Targets..... 109
 - Unexporting NVMe-oF Targets 109
- Using Consistency Checks to Ensure Data Integrity 110
 - Starting Consistency Checks Manually 111
 - Stopping Consistency Check..... 112

| | |
|---|------------|
| Scheduling Consistency Checks..... | 112 |
| Viewing Consistency Check Information | 113 |
| Setting the Consistency Check Policy..... | 114 |
| Excluding Drive Groups from the Consistency Check Policy | 115 |
| ADDITIONAL FUNCTIONS..... | 116 |
| Configuring Boot-Drive Devices | 116 |
| Procedure for CentOS | 117 |
| Procedure for SLES 15 SP2, and SP3 | 125 |
| Manually Migrating the RAID Configuration Between Hosts..... | 127 |
| Restoring a RAID Configuration from a Backup Configuration File..... | 127 |
| Restoring a RAID Configuration from SSD Metadata | 128 |
| Restarting the SupremeRAID™ Service After Upgrading the System Kernel | 129 |
| Obtaining SMART Information from Devices..... | 130 |
| Monitoring System Input/Output Statistics for Devices Using iostat | 133 |
| sysstat Versions v12.3.3 and Later | 133 |
| sysstat Versions Prior to v12.3.3 | 135 |
| Setting Up the Auto-mount File Systems on Linux Using the SupremeRAID™ Driver.. | 136 |
| ESXi Virtual Machine Support Using GPU Passthrough | 139 |
| Configuring Hosts for NVIDIA GPU Device Passthrough | 139 |
| Configuring Virtual Machines | 140 |
| Using Self-Encrypting Drives | 142 |
| Importing a Single SED Key Using NQN/WWID | 142 |
| Importing a Batched SED Key Using NQN/WWID | 143 |
| Displaying SED Key Information..... | 143 |
| Deleting SED Keys..... | 143 |
| TROUBLESHOOTING..... | 145 |
| Sequential Read Performance is Not as Expected on a New Drive Group | 145 |
| Kernel Log Message "failed to set APST feature (-19)" Appears When Creating Physical Drives | 145 |

Decoding LED Patterns on the Backplane..... 145

Received "The arch of the controller and graid software mismatched" Message When
Applying License..... 146

SAFETY INFORMATION 149

English Version 149

Chinese Version..... 150

Chinese Version (TC) 151

INTRODUCTION

SupremeRAID™ is the most powerful, high-speed data protection solution specially designed for NVMe SSDs. SupremeRAID™ installs a virtual NVMe controller onto the operating system and integrates a high-performance, GPU-base PCIe RAID card into the system to manage the RAID operations of the virtual NVMe controller.

This document explains how to install the SupremeRAID™ software package for Linux and how to manage the RAID components using the command-line interface.

Software Module Overview

The SupremeRAID™ Software module has the following major components:

- **graidctl** – command-line management tool
- **graid_server** – management daemon that handles requests from **graidctl** to control the driver
- **graid.ko** – driver kernel module
- **graid_core** – GPU instance


SupremeRAID™ Specification

| SupremeRAID™ Driver Specifications | |
|---|---|
| Supported models: | SR-1000, SR-1010, SR-1001 |
| Supported RAID levels: | RAID 0, 1, 5, 6, 10 |
| Recommended minimum drive number for each RAID level: | RAID 0 : at least two drives RAID 1 : at least two drives RAID 5 : at least three drives RAID 6 : at least four drives RAID 10 : at least four drives |
| Maximum number of physical drives: | 32 |
| Maximum number of drive groups: | 8 |
| Maximum number of virtual drives per drive group: | 1,023 |
| Maximum size of the drive group: | Defined by the physical drive sizes |
| Configurable strip size (RAID0, RAID10) | 4k, 8k, 16k, 32k, 64k,128k |

RAID Components

SupremeRAID™ has three major RAID logical components:

- Physical Drive (PD)
- Drive Group (DG)
- Virtual Drive (VD)



Physical Drive (PD)

Since NVMe drives are not directly attached to the SupremeRAID™ controller, you must tell the controller which SSDs can be managed. After an SSD is created as a physical drive, the SupremeRAID™ driver unbinds the SSD from the operating system, meaning the device node (`/dev/nvmeX`) disappears and is no longer accessible. At the same time, the SupremeRAID™ driver creates a corresponding device node (`/dev/gpdX`). You can check the SSD information, such as SSD model or SMART logs, using this device node. To control and access the SSD using `/dev/nvmeXn1`, you must first delete the corresponding physical drive.

SupremeRAID™ supports 32 physical drives, regardless of whether the physical drives are created from a native NVMe SSD, a drive connected through NVMe-oF, or a SAS/SATA disk.

Drive Group (DG)

The main component of RAID logic is a RAID group. When the drive group is created, the SupremeRAID™ driver initializes the physical drives with the corresponding RAID mode to ensure that the data and parity are synchronized.

There are two types of initialization processes.

- **Fast Initialization:** When all of the physical drives in the drive group (DG) support the de-allocate dataset management command, the SupremeRAID™ driver performs fast initialization by default, which optimizes the drive group state immediately.
- **Background Initialization:** Performance is slightly affected by the initialization traffic, but you can still create the virtual drive and access the virtual drive during a background initialization.

SupremeRAID™ supports eight drive groups, with a maximum of 32 physical drives in one drive group.

Virtual Drive (VD)

The virtual drive is equivalent to the RAID volume. You can create multiple virtual drives in the same drive group for multiple applications. The corresponding device node (`/dev/gdgXnY`) appears on the operating system when you create a virtual drive, and you can make the file system or running application directly on this device node. Currently, the SupremeRAID™ driver supports a maximum of 1023 virtual drives in each drive group.

Note: If you upgrade from version 1.2.x to version 1.5.x of the graid driver, the device path changes from `/dev/gvdXn1` to `/dev/gdgXnY`.

Features Overview

The SupremeRAID™ presents a range of features that facilitate convenient data storage methods and incorporate diverse protection mechanisms to ensure data integrity. The following will outline key features that contribute to achieving our objectives and fostering a foundational understanding of our services.

Ensuring Data Integrity with Consistency Checks

The SupremeRAID™ is designed to provide high reliability and data integrity levels. A key feature that enables this is the consistency check function.

The consistency check function allows administrators to ensure that the data stored on the SupremeRAID™ system is intact and uncorrupted. These checks can be performed on a regular schedule or manually initiated as needed. When a consistency check is completed, the system compares the data on each disk to identify any discrepancies or errors.

Depending on the settings chosen by the administrator, the consistency check function can either automatically fix any errors that are found or stop the check and alert the administrator to any detected errors. This feature provides administrators with flexibility and control over how the system responds to errors.


For detailed information about graid commands for the consistency check, see [Using Consistency Checks to Ensure Data Integrity](#) on page 108.

Note: The consistency check function is not supported on SupremeRAID™ systems configured in RAID0 mode because RAID0 does not provide data redundancy and does not require data consistency checks.

SupremeRAID's Dual-Controller Architecture for Auto-Failover and High-Availability

This feature enables the SupremeRAID™ system to automatically fail over to another SupremeRAID™ card when one SupremeRAID™ card experiences an issue without any interruption in service. This increased reliability and availability ensures that the system remains operational even in the event of a single card failure.

SupremeRAID™ supports dual-controller configurations in two modes: dual-active and active-passive. This enhances our RAID solution with comprehensive protection and security. Additionally, the high availability (HA) functionality remains unaffected by the root complex. Whether within the same root complex or across different root complexes, we have implemented failover mechanisms to ensure data integrity.




Setting Up the NVMe-oF Initiator Server and Managing Your RAID Components

The SupremeRAID™ allows you to easily manage a remote target server or storage pool that uses NVMe-over-Fabrics (NVMe-oF) technology. Both TCP and RDMA connections are supported, providing flexibility and compatibility with a wide range of systems. With the SupremeRAID™, you can create a virtual volume with RAID capabilities without the need for reconfiguration or re-cabling on the host server. This allows you to take advantage of the benefits of NVMe-oF, including increased capacity and improved data protection.

For detailed information about graid commands for the NVMe-oF initiator, see

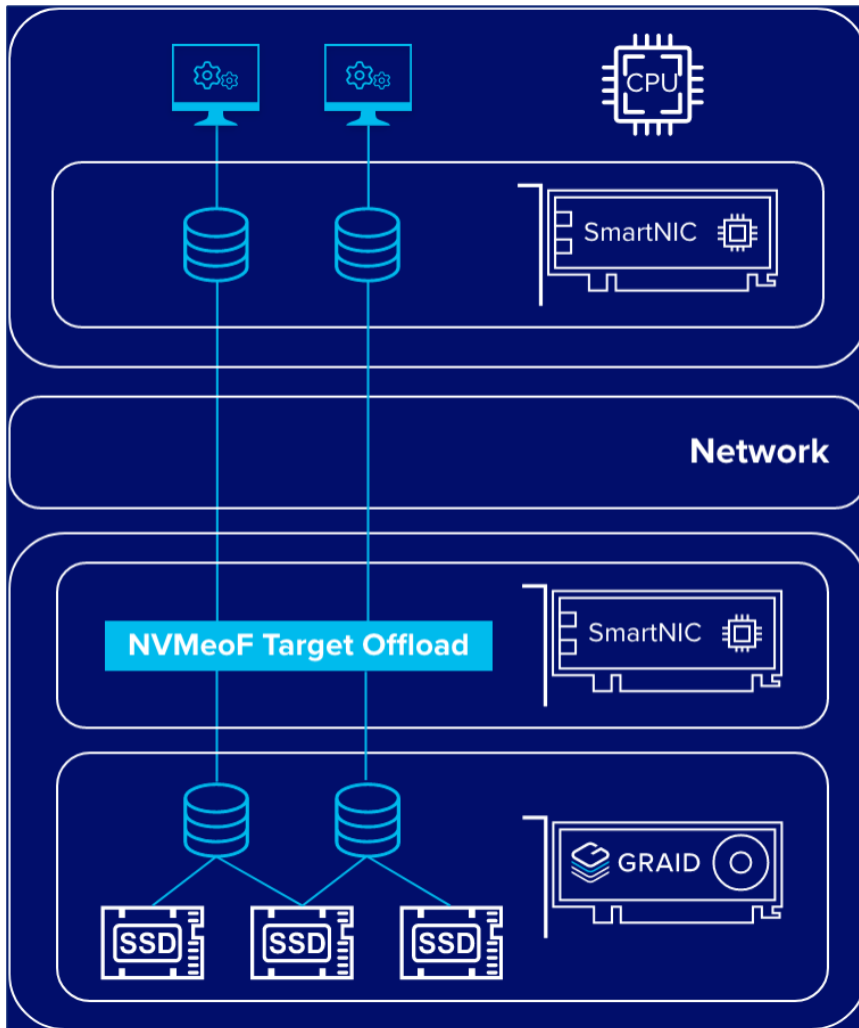
[Managing Remote NVMe-oF Targets](#) on page 103.



Sharing the SupremeRAID™ Volume as a NVMe-oF Target Server

The SupremeRAID™ allows you to easily compose local NVMe devices into a RAID array and share that array as an NVMe- over-Fabrics (NVMe-oF) target server. By using a SmartNIC to accelerate data transfer, you can achieve low latencies and high performance for your remote NVMe-oF clients.

For detailed information about graid commands for the NVMe-oF target, see [Exporting NVMe-oF Target Management](#) on page 104.



SPDK BDEV Feature of SupremeRAID™

The SupremeRAID™ software incorporates SPDK (Storage Performance Development Kit) feature, enabling direct access to operate the NVMe queue from user space through the SupremeRAID™ native BDEV (Block Device) interface. This integration offers significant benefits that enhance the overall performance and efficiency of the system.

The SPDK feature facilitates direct user application access to NVMe queues from user space. This minimizes data access and processing latency, resulting in enhanced system responsiveness through reduced overhead and fewer context switches. Moreover, this direct access eliminates the necessity for data transfers between user space and kernel space, thereby decreasing CPU utilization caused by kernel module activity. This optimization enables the CPU to prioritize crucial tasks, leading to improved overall system performance.

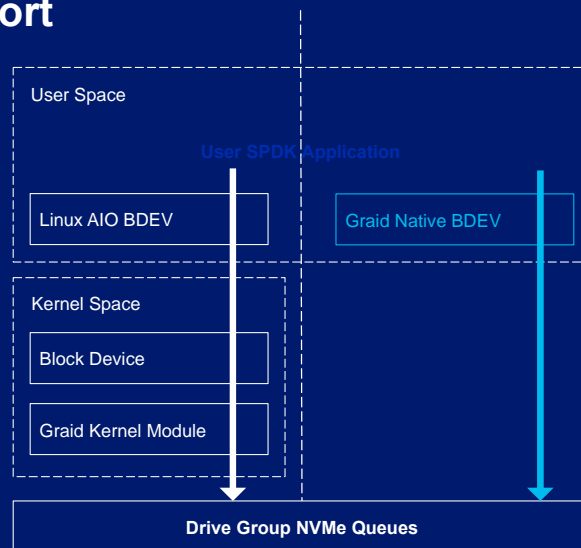
The SPDK feature in SupremeRAID™ contributes to an optimized storage solution, particularly in high-performance scenarios, where latency reduction and improved CPU utilization are crucial factors. By harnessing the power of SPDK, Graid ensures that users can maximize the potential of their NVMe devices while experiencing enhanced data processing capabilities with minimal overhead.

SPDK BDEV Interface Support

- The user application has direct access to operate the NVMe queue from user space via Graid native BDEV interface.

This:

- **Reduces latency**
- **Decreases CPU utilization** consumed by the kernel module



INSTALLATION

This section describes how to install the SupremeRAID™ hardware and software package for Linux operating systems.

Prerequisites

Before proceeding with the installation, make sure the system meets the following requirements:

- Minimum system requirements
 - CPU: 2 GHz or faster with at least 8 cores
 - RAM: 16 GB
 - Supported operating system: see page 16
 - An available PCIe Gen3 or Gen4 x16 slot
- The SupremeRAID™ card must be installed into a PCIe x16 slot.
- The IOMMU function (AMD) or VT-d function (Intel) is usually disabled in the system BIOS, typically found on the BIOS Advanced page.
- It is highly recommended to disable the UEFI Secure Boot function on the BIOS security page, If UEFI Secure Boot is not applicable in your system, you will need to sign the NVIDIA Kernel Module. For further information and troubleshooting, please refer to the Nvidia website.
- The SupremeRAID™ software package, which includes the Pre-Installer and Installer, can be downloaded directly from the Graid Technology website. The Pre-Installer configures all necessary dependencies and environment settings automatically prior to installing the graid driver. The Installer contains the graid driver package and will automatically detect your Linux distributions and install the appropriate files.
- Make sure a SupremeRAID™-compatible SSD drive is being used. For a list of compatible drives, see the Drivers & Documentation section on our website.

Note: To use virtualization services such as ESXi, you must enable the IOMMU (AMD) or VT-d (Intel) function. For more information, see [ESXi Virtual Machine Support Using GPU Passthrough](#) on page 137.

Supported Operating Systems

Graid has been tested with the operating system versions in the following table. For other operating system versions, contact Graid Technology support.

| Linux Distro | X86_64 | Arm64 | Support Kernel Version |
|--------------|--------|--------------------------|------------------------|
| RHEL | 7.9 | Not currently compatible | 3.10.0 |
| | 8.3 | Not currently compatible | 4.18.0 |
| | 8.4 | Not currently compatible | 4.18.0 |
| | 8.5 | Not currently compatible | 4.18.0 |
| | 8.6 | Not currently compatible | 4.18.0 |
| | 8.7 | Not currently compatible | 4.18.0 |
| | 8.8 | Not currently compatible | 4.18.0 |
| | 9.0 | Not currently compatible | 5.14.0 |
| | 9.1 | Not currently compatible | 5.14.0 |
| Rocky Linux | 8.5 | Not currently compatible | 4.18.0 |
| | 8.6 | Not currently compatible | 4.18.0 |
| | 8.7 | Not currently compatible | 4.18.0 |
| | 8.8 | Not currently compatible | 4.18.0 |
| AlmaLinux | 8.5 | Not currently compatible | 4.18.0 |
| | 8.6 | Not currently compatible | 4.18.0 |
| | 8.7 | Not currently compatible | 4.18.0 |
| | 8.8 | Not currently compatible | 4.18.0 |

| Linux Distro | X86_64 | Arm64 | Support Kernel Version |
|---------------|---------|--------------------------|------------------------|
| Ubuntu | 20.04.0 | 20.04.0 | 5.15.0 |
| | 20.04.1 | 20.04.1 | 5.15.0 |
| | 20.04.2 | 20.04.2 | 5.15.0 |
| | 20.04.3 | 20.04.3 | 5.15.0 |
| | 20.04.4 | 20.04.4 | 5.15.0 |
| | 20.04.5 | 20.04.5 | 5.15.0 |
| | 22.04.2 | 22.04.0 | 5.15.0 |
| openSUSE Leap | 15.2 | Not currently compatible | 5.3.18 |
| | 15.3 | Not currently compatible | 5.3.18 |
| SLES | 15 SP2 | Not currently compatible | 5.3.18 |
| | 15 SP3 | Not currently compatible | 5.3.18 |
| CentOS | 7.9 | Not currently compatible | 3.10.0 |
| | 7.9 | Not currently compatible | 4.18.0 |
| | 8.3 | Not currently compatible | 4.18.0 |
| | 8.4 | Not currently compatible | 4.18.0 |
| | 8.5 | Not currently compatible | 4.18.0 |

Note: Our product has been rigorously tested for compatibility with specific operating system versions. If you use an operating system version other than one mentioned in this user guide, we recommend you contact our support team to determine the level of support we can offer. Installing the Ubuntu Server version is considered best practice. Installing the Ubuntu Desktop version may result in kernel compatibility issues with the SupremeRAID™ driver.

Tested NVMe Devices

The following NVMe drives passed Graid Technology qualification and can be used with SupremeRAID™. Graid Technology updates this list when new NVMe drives pass the qualification process. For the latest information, see the Compatible NVMe Drives List on the Graid Technology website.

| Manufacturer | Series | Interface | Form Factor |
|-------------------|----------------|--------------|--------------|
| Dapustor | R5100 | PCIe Gen 4x4 | 2.5 inch U.2 |
| Dapustor | R5103 | PCIe Gen 4x4 | 2.5 inch U.2 |
| Hagiwara Solution | JN2E | PCIe Gen 4x4 | 2.5 inch U.2 |
| Solidigm | DC P4510 | PCIe Gen 3x4 | 2.5 inch U.2 |
| Solidigm | DC P4610 | PCIe Gen 3x4 | 2.5 inch U.2 |
| Solidigm | D5-P5316 | PCIe Gen 4x4 | 2.5 inch U.2 |
| Solidigm | D7-P5510 | PCIe Gen 4x4 | 2.5 inch U.2 |
| Solidigm | D7-P5520 | PCIe Gen 4x4 | 2.5 inch U.2 |
| Solidigm | D7-P5620 | PCIe Gen 4x4 | 2.5 inch U.2 |
| Solidigm | Optane™ P5800X | PCIe Gen 4x4 | 2.5 inch U.2 |
| Kingston | DC1500M | PCIe Gen 3x4 | 2.5 inch U.2 |
| Kioxia | CD5 | PCIe Gen 3x4 | 2.5 inch U.2 |
| Kioxia | CD6 | PCIe Gen 4x4 | 2.5 inch U.2 |
| Kioxia | CM6 | PCIe Gen 4x4 | 2.5 inch U.2 |
| Kioxia | CD8 | PCIe Gen 4x4 | 2.5 inch U.2 |
| Kioxia | CM7 | PCIe Gen 5x4 | 2.5 inch U.2 |
| Memblaze | P6536 | PCIe Gen 4x4 | 2.5 inch U.2 |
| Micron | 7300 PRO | PCIe Gen 4x4 | 2.5 inch U.2 |
| Micron | 7450 | PCIe Gen 4x4 | 2.5 inch U.2 |
| Phison | EPW5900 | PCIe Gen 4x4 | 2.5 inch U.2 |
| Samsung | PM983 | PCIe Gen 3x4 | 2.5 inch U.2 |

| Manufacturer | Series | Interface | Form Factor |
|-----------------|-------------|--------------|--------------|
| Samsung | PM9A3 | PCIe Gen 4x4 | 2.5 inch U.2 |
| Samsung | PM1733 | PCIe Gen 4x4 | 2.5 inch U.2 |
| Samsung | PM1743 | PCIe Gen 4x4 | 2.5 inch U.2 |
| ScaleFlux | CSD-3000 | PCIe Gen 4x4 | 2.5 inch U.2 |
| Seagate | NYTRO 5550H | PCIe Gen 4x4 | 2.5 inch U.2 |
| Western Digital | SN640 | PCIe Gen 3x4 | 2.5 inch U.2 |
| Western Digital | SN650 | PCIe Gen 3x4 | 2.5 inch U.2 |

Installing the Hardware

ESD Warning


Electronic components and circuits are sensitive to ElectroStatic Discharge (ESD). When handling any circuit board assemblies including Connect Tech carrier assemblies, it is recommended that ESD safety precautions be observed. ESD safe best practices include, but are not limited to:

- Leaving circuit boards in their antistatic packaging until they are ready to be installed.
- Using a grounded wrist strap when handling circuit boards, at a minimum you should touch a grounded metal object to dissipate any static charge that may be present on you.
- Only handling circuit boards in ESD safe areas, which may include ESD floor and table mats, wrist strap stations and ESD safe lab coats.
- Avoiding handling circuit boards in carpeted areas.
- Try to handle the board by the edges, avoiding contact with components.

Installation Procedure


Perform the following procedure to install SupremeRAID™ into your system.

- Step 1** Power down your system.
- Step 2** Unplug the power cord from the AC power source.
- Step 3** Remove the side panel from your system to gain access to the motherboard.
- Step 4** If your system has a PCIe card, remove it. If a retention bar is holding the card in place, remove the screw securing the card. If there is no existing PCIe card, remove the access covers from the primary x16 PCI express slot.



Note: The SupremeRAID™ SR-1010 is dual-slot card and requires you to remove two adjacent slot covers. The SupremeRAID™ SR-1000 and SupremeRAID™ SR-1001 are single slot cards and require only a single- slot.

Step 5 Install the card into the primary x16 PCI Express slot. Press gently on the card until it is seated securely in the slot and reattach the SupremeRAID™ card bracket retention mechanism.



Note: Install the SupremeRAID™ card into the primary x16 PCI Express slot. The SupremeRAID™ SR-1010 is dual-slot card and covers the adjacent slot. The SupremeRAID™ SR-1000 and SupremeRAID™ SR-1001 are single-slot cards. For more information, see <https://manuals.plus/nvidia/rtx-ampere-architecture-based-graphics-card-manual#ixzz7wk7PysLh>.

Step 6 Secure the card to the system frame using the screw(s) you removed in step 4.

Step 7 Install the side panel you removed in step 3.

Installing the Software Driver

The recommended and quickest way to install the graid software is by using the pre-installer scripts and installer (described below).

However, if you prefer to install the software manually or your environment lacks Internet access, follow the procedure on page 30 to configure the environment settings and install the Graid driver manually. If you have already installed the software and only wish to upgrade it, please refer to the instructions on the page 53 for the upgrade configuration.

Using the Pre-installer and Installer

The graid pre-installer is an executable file that contains the required dependencies and a setup script that installs the NVIDIA driver. The script makes it easy to prepare the environment and install the SupremeRAID™ driver in every supported Linux distribution. Use the following steps to prepare the environment and install the SupremeRAID™ driver using the pre-installer in supported Linux distributions.


Note: To run the pre-installer, the system must have internet access to download the required dependencies from the official mirror.

Step 1 Go to the Graid Technology website to download the latest version of the pre-installer and make it executable.

Drivers & Documentation

```
$ sudo chmod +x <filename>
```

Dependencies and Utilities

| | Links |
|----------------------------|---|
| NVIDIA Driver | NVIDIA-Linux-x86_64-515.86.01.run |
| SupremeRAID™ Pre-installer | graid-sr-pre-installer-1.5.0-run |

Step 2 Execute the pre-installer and follow the instructions to complete the pre-installation process, as shown in the following figure.

```

root@graid-demo:~# ./graid-sr-pre-installer-1.5.0-55.run
Reading package lists... Done
Building dependency tree
Reading state information... Done
gawk is already the newest version (1:5.0.1+dfsg-1).
pciutils is already the newest version (1:3.6.4-1ubuntu0.20.04.1).
tar is already the newest version (1.30+dfsg-7ubuntu0.20.04.3).
The following packages will be upgraded:
  mokutil
1 upgraded, 0 newly installed, 0 to remove and 259 not upgraded.
Need to get 26.6 kB of archives.
After this operation, 13,3 kB of additional disk space will be used.
Get:1 http://tw.archive.ubuntu.com/ubuntu focal-updates/main amd64 mokutil amd64 0.6.0-2-20.04.1 [26.6 kB]
Fetched 26.6 kB in 1s (20.8 kB/s)
(Reading database ... 192964 files and directories currently installed.)
Preparing to unpack .../mokutil_0.6.0-2-20.04.1_amd64.deb ...
Unpacking mokutil (0.6.0-2-20.04.1) over (0.3.0+1538710437.fb6250f-1) ...
Setting up mokutil (0.6.0-2-20.04.1) ...
Processing triggers for man-db (2.9.1-1) ...
Extracting installer files, please wait a few seconds ...
Extracting installer files done.

Starting graid-preinstaller ...
This graid-preinstaller is for Ubuntu
Running Graid service check
Graid service check succeeded.

Prepare install NVIDIA driver
Checking Xorg ...
Checking nouveau ...
Running install NVIDIA Driver. (This step will take a while.)
Mon Aug 28 19:03:00 2023
-----+-----
| NVIDIA-SMI 515.86.01    Driver Version: 515.86.01    CUDA Version: 11.7    |
|-----+-----|
| GPU Name      Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
|               |                 |          |             MIG M. |
|-----+-----|
| 0   NVIDIA RTX A2000    Off   | 00000000:22:00:0 Off |             0         |
| 30%   50C   P2      24W / 70W | 0MiB / 5754MiB |      2%   Default |
|               |                 |          |             N/A     |
|-----+-----|
+-----+-----+
| Processes: |
| GPU  GI  CI           PID  Type  Process name                        GPU Memory |
| ID   ID  ID           |                   |           | Usage |
|-----+-----|
| No running processes found |
+-----+-----+
Install NVIDIA Driver succeeded.

This graid-preinstaller will reboot the system for apply previous setting!
Do you want to continue? [Y/n]

```




Step 3 After running the pre-installation script, type **Y** when prompted to reboot the system.

Step 4 Go to the Graid Technology website, download the latest version of the installer, and make it executable.

Drivers & Documentation

```
$ sudo chmod +x <filename>
```

Driver Packages

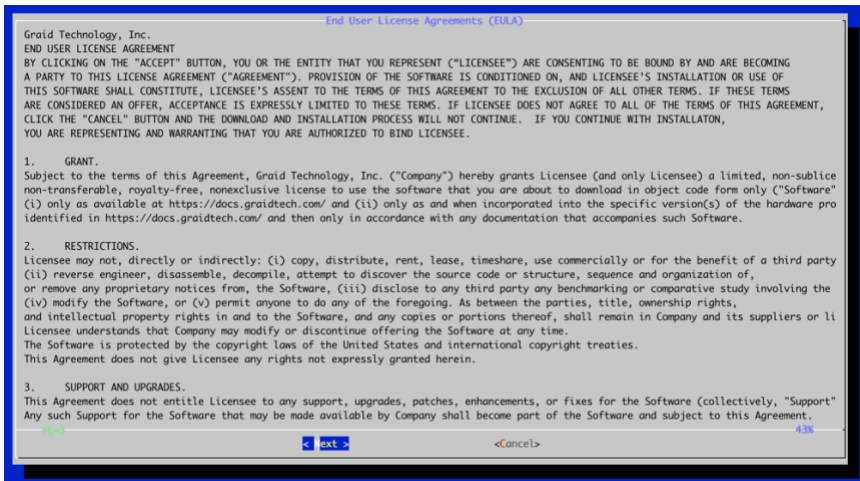
| Model | x86_64 |
|---------|--|
| SR-1000 | graid-sr-installer-1.5.0-.run |
| SR-1001 | graid-sr-installer-1.5.0-.run |
| SR-1010 | graid-sr-installer-1.5.0-.run |

Step 5 Execute the installer and follow the provided steps to complete the installation.

A At the Welcome page, select **Next** and click **Enter** to view the end-user license agreement.




B In the end-user license agreement page, use the spacebar to scroll down the content. After you review the license, select **Next** and click **Enter**.




C Type **accept**, click **tab**, select **next**, and click **enter** to accept the license agreement.



D Confirm the installation package, and then Click **Next** to continue with the installation.



E Complete the installation, and the installer will reboot system.



Step 6 To activate the software, apply the SupremeRAID™ license key.

```
$ sudo graidctl apply license <LICENSE_KEY>
```

Using Installer for Silent Installation

This section is designed for users who require mass deployment and may be designing scripts for installation. However, we strongly recommend using the GUI installation process for the best user experience and comprehensive configuration options.

Step 1 Please follow the steps from the previous section to download the pre-installer and installer, make them executable, and use the pre-installer to install the dependencies required by the Graid service.

```
$ sudo chmod +x <filename>
```

Step 2 To install the driver directly with the license acceptance, simply add the command '--accept-license' at the end when executing the installer.

```
$ sudo ./<filename> --accept-license
```

```
root@graid-demo:~/driver# sudo ./graid-sr-installer-1.5.0-001-659-82-x86_64.run --accept-license
Extracting installer files, please wait a few seconds ...
Extracting installer files done.

Starting installer ...
systemctl stop graid
Creating symlink /var/lib/dkms/grebar/0.1.0/source -> /usr/src/grebar-0.1.0

Kernel preparation unnecessary for this kernel. Skipping...

Building module:
cleaning build area...
make -j32 KERNELRELEASE=5.15.0-83-generic...
Signing module:
- /var/lib/dkms/grebar/0.1.0/5.15.0-83-generic/x86_64/module/grebar.ko
Secure Boot not enabled on this system.
cleaning build area...

grebar.ko:
Running module version sanity check.
- Original module
- No original module exists within this kernel
- Installation
- Installing to /lib/modules/5.15.0-83-generic/updates/dkms/

depmod...
Selecting previously unselected package graid-sr.
(Reading database ... 83819 files and directories currently installed.)
Preparing to unpack ../graid-sr-1.5.0-659.g10e76f72.001.x86_64.deb ...
No need to patch
Unpacking graid-sr (1.5.0) ...
Setting up graid-sr (1.5.0) ...
Creating symlink /var/lib/dkms/graid/1.5.0/source -> /usr/src/graid-1.5.0

Kernel preparation unnecessary for this kernel. Skipping...

Building module:
cleaning build area...
./build.sh 5.15.0-83-generic /lib/modules/5.15.0-83-generic/build.....
Signing module:
- /var/lib/dkms/graid/1.5.0/5.15.0-83-generic/x86_64/module/graid-nvidia.ko
- /var/lib/dkms/graid/1.5.0/5.15.0-83-generic/x86_64/module/graid.ko
Secure Boot not enabled on this system.
cleaning build area...

graid.ko:
Running module version sanity check.
- Original module
- No original module exists within this kernel
- Installation
- Installing to /lib/modules/5.15.0-83-generic/updates/dkms/

graid-nvidia.ko:
Running module version sanity check.
- No original module exists within this kernel
- Installation
- Installing to /lib/modules/5.15.0-83-generic/updates/dkms/

depmod...
Processing triggers for man-db (2.10.2-1) ...
Suggestion!! This installer will reboot the system for apply previous kernel module grebar setting!
Do you want to continue? [Y/n] █
```

Step 3 To activate the software, apply the SupremeRAID™ license key.

```
$ sudo graidctl apply license <LICENSE_KEY>
```

Manual Installation

The following procedure describes how to manually install the Graid software on various operating systems. The reference for packages and dependencies for each operating system is provided below.

- For CentOS, Rocky Linux, AlmaLinux, and RHEL operating systems, see page 31.
- For Ubuntu operating systems see page 34.
- For openSUSE operating systems, see page 36.
- For SLES operating systems, see page 38.

Note: If the system does not have internet access, download the required dependencies from the official repositories. For a detailed description of distributions, see the distribution section below. Do not perform this procedure unless you need to install dependencies manually or the pre-installer procedure did not work for you. Otherwise, see Supported Operating Systems on page 16 and use the automated pre-installer script to install the graid software

Dependency Table for Manual Installation

Here is the dependency tree for manual installation and the comparison table for each operating system.

| RHEL | CentOS/Rocky/ Almalinux/Oracle | SLES | Debian/Ubuntu |
|----------|-----------------------------------|----------|---------------|
| automake | automake | automake | automake |
| dialog | dialog | dialog | dialog |
| dkms | dkms | dkms | dkms |
| gcc | gcc | gcc | gcc |
| ipmitool | ipmitool | ipmitool | ipmitool |
| make | make | make | make |
| mdadm | mdadm | mdadm | mdadm |
| mokutil | mokutil | mokutil | mokutil |
| pciutils | pciutils | pciutils | pciutils |
| tar | tar | tar | tar |

| | | | |
|-----------------------------------|-----------------------------------|---|----------------------------------|
| vim | vim | vim | vim |
| wget | wget | wget | wget |
| sg3_utils | sg3_utils | libsgutils-devel | libsgutils2-2 |
| -- | -- | libpci3 | libpci3 |
| -- | -- | libpci3 | libpci3 |
| sqlite-libs | sqlite-libs | sqlite3 | sqlite3 |
| -- | -- | libudev-devel | -- |
| -- | -- | -- | initramfs-tools |
| -- | -- | -- | gawk |
| gcc-c++-\${VERSION_ID} | gcc-c++ | g++ | g++ |
| gcc-\${VERSION_ID} | -- | -- | -- |
| kernel-devel-\${kernel_version} | kernel-devel-\${kernel_version} | -- | -- |
| kernel-headers-\${kernel_version} | kernel-headers-\${kernel_version} | -C kernel-default-devel=\${kernel_version}_s use) | linux-headers-\${kernel_version} |

Note: To determine the kernel version for RHEL, you can use the command `uname -r`. For SUSE, extract the kernel version using `uname -r | awk -F"-default" '{print $1}'`. Additionally, please using `awk -F=' /VERSION_ID/{ gsub("/", ""); print $2}' /etc/os-release` to retrieve the version ID.

Manual Installation on a CentOS, Rocky Linux, AlmaLinux, and RHEL Operating Systems

Graid Technology, Inc. recommends referring to Supported Operating Systems on page 19 and using the pre-installer to configure the environmental settings.

Step 1 Install the package dependencies and build for Dynamic Kernel Module Support (DKMS) based on your operating system.

- For CentOS, Rocky Linux, and AlmaLinux: issue the following commands.

```
$ sudo yum install --enablerepo=extras epel-release
$ sudo yum install vim wget make automake gcc gcc-c++ kernel-devel
kernel-headers kernel dkms ipmitool tar mdadm sg3_utils sqlite-libs
automake dialog
```

- For RHEL8, issue the following commands:

```
$ sudo yum install https://dl.fedoraproject.org/pub/epel/epel-release-
latest-8.noarch.rpm
$ sudo yum install vim wget make automake kernel-devel-$(uname -r)
kernel-headers-$(uname -r) dkms gcc gcc-c++ ipmitool tar mdadm
sg3_utils sqlite-libs automake dialog
```

- For RHEL7.9: issue the following commands.

```
$ sudo yum install https://dl.fedoraproject.org/pub/epel/epel-release-
latest-7.noarch.rpm
$ sudo yum install gcc-$(awk -F=' ' /VERSION_ID/{ gsub("/", ""); print
$2}' /etc/os-release) gcc-c+-$(awk -F=' ' /VERSION_ID/{ gsub("/", "");
print $2}' /etc/os-release)
$ sudo yum install vim wget make automake kernel-devel-$(uname -r)
kernel-headers-$(uname -r) dkms
ipmitool tar mdadm sg3_utils sqlite-libs automake dialog
```

Step 2 Add the kernel option. This step prevents the Nouveau driver from loading during installation and disables IOMMU in the system BIOS.

```
$ sudo vim /etc/default/grub
```

Step 3 Append the command line parameters and then update the grub configuration based on your operating system.

- For RHEL8, append `iommu=pt` and `nvme_core.multipath=Y` to `GRUB_CMDLINE_LINUX_DEFAULT`.
- For RHEL7.9, append `iommu=pt` to `GRUB_CMDLINE_LINUX_DEFAULT`.

```
$ sudo grub2-mkconfig -o /boot/grub2/grub.cfg
```

Step 4 Append `blacklist nouveau` and `options nouveau modeset=0` to the end of the `/etc/modprobe.d/graid-blacklist.conf` file to disable the Nouveau driver and update `initramfs`.

```
root@graid-demo:/etc/modprobe.d# cat graid-blacklist.conf
blacklist nouveau
options nouveau modeset=0
```

Note: You might need to manually create the `/etc/modprobe.d/graid-blacklist.conf` file and append `blacklist nouveau` and `options nouveau modeset=0`.

```
$ sudo update-initramfs -u
```

- For CentOS, Rocky Linux, and AlmaLinux: Find the latest version of the kernel and assign it to `-kver`.

```
$ sudo dracut -f --kver `rpm -qa | grep kernel-headers | awk -F'kernel-headers-' {'print $2}'`
```

- For RHEL: issue the following command.

```
$ sudo dracut -f
```

Step 5 Reboot the system and make sure the grub configuration was applied. You can check `/proc/cmdline` for the grub configuration in use. For example:

- For RHEL8:

```
[root@graid ~]# cat /proc/cmdline
BOOT_IMAGE=(hd3,gpt8)/vmlinuz-4.18.0-305.17.1.el8_4.x86_64 root=UUID=ba33c54d-74c9-409d-ae05-db27cacd68b3 ro crashkernel=auto resume=UUID=ae5b3808-b657-4593-a598-c5cbc5a87105 rhgb quiet rd.driver.blacklist=nouveau iommu=pt nvme_core.multipath=Y
```

- For RHEL7:

```
[root@graid ~]# cat /proc/cmdline
BOOT_IMAGE=(hd3,gpt8)/vmlinuz-4.18.0-305.17.1.el8_4.x86_64 root=UUID=ba33c54d-74c9-409d-ae05-db27cacd68b3 ro crashkernel=auto resume=UUID=ae5b3808-b657-4593-a598-c5cbc5a87105 rhgb quiet rd.driver.blacklist=nouveau iommu=pt
```

Step 6 Install the NVIDIA driver.

```
$ wget https://tw.download.nvidia.com/XFree86/Linux-
x86_64/515.86.01/NVIDIA-Linux-x86_64-515.86.01.run
$ chmod +x ./NVIDIA-Linux-x86_64-515.86.01.run
```

- For CentOS: Use the latest version of kernel-headers to install the NVIDIA driver.

```
$ sudo ./NVIDIA-Linux-x86_64-515.86.01.run -s --no-systemd --no-opengl-
files --no-nvidia-modprobe -- dkms -k `rpm -qa | grep kernel-headers |
awk -F'kernel-headers-' {'print $2'}`
```

- For RHEL:

```
$ sudo ./NVIDIA-Linux-x86_64-515.86.01.run -s --no-systemd --no-opengl-
files --no-nvidia-modprobe -- dkms -k `rpm -qa | grep kernel-headers |
awk -F'kernel-headers-' {'print $2'}`
```

Note: The Nouveau driver is now disabled. Reboot and install the NVIDIA driver before proceeding with the installation.

Step 7 Use the `nvidia-smi` command to confirm that the NVIDIA GPU is working. The following figure shows an output example of a successful installation.

```
root@graid-demo:~# nvidia-smi
Tue Aug 29 18:18:37 2023

+-----+
| NVIDIA-SMI 515.86.01   Driver Version: 515.86.01   CUDA Version: 11.7   |
+-----+-----+
| GPU   Name           Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|  Memory-Usage | GPU-Util  Compute M. |
|                                           MIG M. |
+-----+-----+
|  0  NVIDIA RTX A2000   Off       | 00000000:22:00.0 Off |          0           |
| 30%  48C   P2      24W / 70W     |  0MiB / 5754MiB |    2%    Default   |
|                                           |                 N/A   |
+-----+-----+

Processes:
+-----+-----+
| GPU   GI   CI          PID   Type   Process name          GPU Memory |
| ID   ID   ID                   |                | Usage          |
+-----+-----+
| No running processes found |
+-----+-----+
```

Step 8 From the Graid Technology website, download the latest version of the installer and make it executable.

```
$ sudo chmod +x <filename>
```

Step 9 Proceed to [Executing the Installer and Completing the Installation](#) on page 43 to execute the installer and to complete the installation.

Manual Installation on an Ubuntu Operating System

Graid Technology, Inc. recommends referring to Supported Operating Systems on page 16 and using the pre-installer to configure the environmental settings.

Step 1 Install the package dependencies and build for DKMS.

```
$ sudo apt-get update

$ sudo apt-get install make automake gcc g++ linux-headers-$(uname -r)
dkms ipmitool initramfs-tools tar mdadm libsgutils2-2 libudev-dev libpci3
sqlite automake dialog
```

Step 2 Disable Ubuntu daily upgrade.

```
$ sed -i '/Unattended-Upgrade "1"/ s/"1"/"0"/' /etc/apt/apt.conf.d/20auto-
upgrades

$ sed -i '/Update-Package-Lists "1"/ s/"1"/"0"/'
/etc/apt/apt.conf.d/20auto-upgrades
```

Step 3 Add the kernel option. This step prevents the Nouveau driver from loading during installation and disables IOMMU in the system BIOS.

```
$ sudo vim /etc/default/grub
```

Step 4 Append `iommu=pt` and `nvme_core.multipath=Y` to `GRUB_CMDLINE_LINUX_DEFAULT`, and then update the grub configuration.

```
$ sudo update-grub
```

Step 5 Append `blacklist nouveau` and `options nouveau modeset=0` to the end of the `/etc/modprobe.d/graid-blacklist.conf` file to disable the Nouveau driver and update `initramfs`.

```
root@graid-demo:/etc/modprobe.d# cat graid-blacklist.conf
blacklist nouveau
options nouveau modeset=0
```

Note: You might need to manually create the `/etc/modprobe.d/graid-blacklist.conf` file and append `blacklist nouveau` and `options nouveau modeset=0`.

```
$ sudo update-initramfs -u
```

Step 6 Reboot the system and make sure the grub configuration was applied. You can check `/proc/cmdline` for the grub configuration in use. For example:

```
root@graid-demo:/etc/modprobe.d# cat /proc/cmdline
BOOT_IMAGE=/boot/vmlinuz-5.15.0-46-generic root=UUID=32b02b62-7173-4f3b-a723-8aa1e2fbf60a ro text iommu=pt nvme_core.multipath=Y
```

Step 7 Install the NVIDIA driver.

```
$ wget https://tw.download.nvidia.com/XFree86/Linux-x86_64/515.86.01/NVIDIA-Linux-x86_64-515.86.01.run
$ sudo chmod +x ./NVIDIA-Linux-x86_64-515.86.01.run
$ sudo ./NVIDIA-Linux-x86_64-515.86.01.run -s --no-systemd --no-opengl-files --no-nvidia-modprobe -- dkms --disable-nouveau
$ sudo reboot
```

Note: The Nouveau driver is now disabled. Reboot and install the NVIDIA driver before proceeding with the installation.

Step 8 Use the `nvidia-smi` command to confirm that the NVIDIA GPU is working. The following figure shows an output example of a successful installation.

```
root@graid-demo:~# nvidia-smi
Tue Aug 29 18:18:37 2023
+-----+
| NVIDIA-SMI 515.86.01      Driver Version: 515.86.01      CUDA Version: 11.7      |
+-----+-----+-----+-----+-----+-----+
| GPU  Name            Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp   Perf   Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
|                                           MIG M.         |
+-----+-----+-----+-----+-----+-----+
|   0   NVIDIA RTX A2000      Off | 00000000:22:00.0 Off |          0          |
| 30%   48C    P2     24W /  70W |  0MiB / 5754MiB |      2%    Default  |
|                                           N/A           |
+-----+-----+-----+-----+-----+
+-----+
| Processes:                                                       GPU Memory |
|  GPU   GI    CI          PID    Type   Process name          Usage   |
|-----+-----+-----+-----+-----+
| No running processes found                                     |
+-----+
```

Step 9 Download the latest version of the SupremeRAID™ driver.

Step 10 From the Graid Technology website, download the latest version of the installer and make it executable.

Drivers & Documentation

```
$ sudo chmod +x <filename>
```

Driver Packages

| Model | x86_64 |
|---------|---|
| SR-1000 | graid-sr-installer-1.5.0-XXXXXXXXXX.run |
| SR-1001 | graid-sr-installer-1.5.0-XXXXXXXXXX.run |
| SR-1010 | graid-sr-installer-1.5.0-XXXXXXXXXX.run |

Step 11 Proceed to [Executing the Installer and Completing the Installation](#) on page 42 to execute the installer and to complete the installation.

Manual Installation on an openSUSE Operating System

Graid Technology, Inc. recommends referring to Supported Operating Systems on page 16 and using the pre-installer to configure the environmental settings.

Step 1 Install openSUSE and select all online repositories.

Step 2 Install the package dependencies and build for DKMS.

```
$ sudo zypper addrepo -f
https://download.opensuse.org/distribution/leap/15.3/repo/oss/ leap-15.3
$ sudo zypper --gpg-auto-import-keys refresh
$ sudo zypper install sudo vim wget libpci3 dkms ipmitool tar mdadm
libsgutils-devel libudev-devel sqlite3 automake dialog
$ sudo zypper install -C kernel-default-devel=$(uname -r | awk -F"-default"
'{print $1}')
```

Step 3 Add the kernel option. This step prevents the Nouveau driver from loading during installation and disables IOMMU in the system BIOS.

```
$ sudo vim /etc/default/grub
```

Step 4 Append `iommu=pt` and `nvme_core.multipath=Y` to `GRUB_CMDLINE_LINUX_DEFAULT`, and then update the grub configuration.

```
$ sudo grub2-mkconfig -o /boot/grub2/grub.cfg
```

Step 5 Append **blacklist nouveau** to the end of the `/etc/modprobe.d/graid-blacklist.conf` file to disable the Nouveau driver.

```
root@graid-demo:/etc/modprobe.d# cat graid-blacklist.conf
blacklist nouveau
options nouveau modeset=0
```

Note: You might need to manually create the `/etc/modprobe.d/graid-blacklist.conf` file and append **blacklist nouveau** and `options nouveau modeset=0`.

Step 6 Set the `allow_unsupported_modules` option to `1` in the `/etc/modprobe.d/10-unsupported-modules.conf` file and update `initrd`.

```
$ sudo mkinitrd
```

Step 7 Reboot the system and make sure the grub configuration was applied. You can check `/proc/cmdline` for the grub configuration in use. For example:

```
root@graid:~# cat /proc/cmdline
BOOT_IMAGE=/boot/vmlinuz-5.3.18-59.5-default root=UUID=7560fe42-0275-4618-b8a0-0785765610c9 modprobe.blacklist=nouveau lommu=pt splash=silent quiet
mitigations=auto nvme_core.multipath=Y
```

Step 8 Install the NVIDIA driver.

```
$ wget https://tw.download.nvidia.com/XFree86/Linux-x86_64/515.86.01/NVIDIA-Linux-x86_64-515.86.01.run
$ sudo chmod +x ./NVIDIA-Linux-x86_64-515.86.01.run
$ sudo ./NVIDIA-Linux-x86_64-515.86.01.run -s --no-systemd --no-opengl-files --no-nvidia-modprobe -- dkms --disable-nouveau
$ sudo reboot
```

Note: The Nouveau driver is now disabled. Reboot and install the NVIDIA driver before proceeding with the installation.

Step 9 Use the `nvidia-smi` command to confirm that the NVIDIA GPU is working. The following figure shows an output example of a successful installation.

```

root@graid-demo:~# nvidia-smi
Tue Aug 29 18:18:37 2023
+-----+
| NVIDIA-SMI 515.86.01    Driver Version: 515.86.01    CUDA Version: 11.7    |
+-----+-----+
| GPU  Name           Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|  Memory-Usage | GPU-Util  Compute M. |
|                                           |              |                  MIG M. |
+-----+-----+
|  0  NVIDIA RTX A2000   Off          | 00000000:22:00:0 Off |      0          0     |
| 30%  48C   P2     24W / 70W |  0MiB / 5754MiB |      2%      Default |
|                                           |              |                  N/A   |
+-----+-----+

+-----+
| Processes:                                                       GPU Memory |
|  GPU   GI    CI          PID    Type   Process name                      Usage    |
|-----+-----+
| No running processes found
+-----+




```

Step 10 From the Graid Technology website, download the latest version of the installer and make it executable.

Drivers & Documentation

```
$ sudo chmod +x <filename>
```

Driver Packages

| Model | x86_64 |
|---------|--|
| SR-1000 | graid-sr-installer-1.5.0-.run |
| SR-1001 | graid-sr-installer-1.5.0-.run |
| SR-1010 | graid-sr-installer-1.5.0-.run |

Step 11 Proceed to [Executing the Installer and Completing the Installation](#) on page 43 to execute the installer and to complete the installation.

Manual Installation on a SLES Operating System

Graid Technology, Inc. recommends referring to Supported Operating Systems on page 16 and using the pre-installer to configure the environmental settings.

Step 1 Install SLES with the following extensions and modules:

- SUSE Package Hub 15 SP3 x86_64
- Desktop Applications Module 15 SP3 x86_64
- Development Tools Module 15 SP3 x86_64

Step 2 Install the package dependencies and build for DKMS.

```
$ sudo zypper addrepo -f
https://download.opensuse.org/distribution/leap/15.3/repo/oss/ leap-15.3

$ sudo zypper --gpg-auto-import-keys refresh

$ sudo zypper install sudo vim wget libpci3 dkms ipmitool tar mdadm
libsgutils-devel libudev-devel sqlite3 automake dialog

$ sudo zypper install -C kernel-default-devel=$(uname -r | awk -F"-
default" '{print $1}')
```

Step 3 Add the kernel option. This step prevents the Nouveau driver from loading during installation and disables IOMMU in the system BIOS.

```
$ sudo vim /etc/default/grub
```

Step 4 Append `iommu=pt` and `nvme_core.multipath=Y` to `GRUB_CMDLINE_LINUX_DEFAULT`, and then update the grub configuration:

```
$ sudo grub2-mkconfig -o /boot/grub2/grub.cfg
```

Step 5 Append `blacklist nouveau` to the end of the `/etc/modprobe.d/graid-blacklist.conf` file to disable the Nouveau driver.

```
root@graid-demo:/etc/modprobe.d# cat graid-blacklist.conf
blacklist nouveau
options nouveau modeset=0
```

Note: You might need to manually create the `/etc/modprobe.d/graid-blacklist.conf` file and append `blacklist nouveau` and `options nouveau modeset=0`.

Step 6 Set the `allow_unsupported_modules` option to `1` in the `/etc/modprobe.d/10-unsupported-modules.conf` file and update `initrd`.

```
$ sudo mkinitrd
```

Step 7 Reboot the system and make sure the grub configuration was applied. You can check `/proc/cmdline` for the grub configuration in use. For example:

```
root@graid:~# cat /proc/cmdline
BOOT_IMAGE=/boot/vmlinuz-5.3.18-59.5-default root=UUID=7560fe42-0275-4618-b8a0-0785765610c9 modprobe.blacklist=nouveau iommu=pt splash=silent quiet
mitigations=auto nvme_core.multipath=Y
```

Step 8 Install the NVIDIA driver.

```
$ wget https://tw.download.nvidia.com/XFree86/Linux-
x86_64/515.86.01/NVIDIA-Linux-x86_64-515.86.01.run
```

```
$ sudo chmod +x ./NVIDIA-Linux-x86_64-515.86.01.run
$ sudo ./NVIDIA-Linux-x86_64-515.86.01.run -s --no-systemd --no-opengl-
files --no-nvidia-modprobe -- dkms --disable-nouveau
$ sudo reboot
```

Note: The Nouveau driver is now disabled. Reboot and install the NVIDIA driver before proceeding with the installation.

Step 9 Use the `nvidia-smi` command to confirm that the NVIDIA GPU is working. The following figure shows an output example of a successful installation.

```
root@graid-demo:~# nvidia-smi
Tue Aug 29 18:18:37 2023
+-----+
| NVIDIA-SMI 515.86.01    Driver Version: 515.86.01    CUDA Version: 11.7    |
+-----+-----+
| GPU  Name           Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp   Perf   Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
|                                           MIG M. |
+-----+-----+
|  0  NVIDIA RTX A2000    Off      | 00000000:22:00.0 Off |          0          |
| 30%   48C   P2     24W / 70W |  0MiB / 5754MiB |    2%      Default  |
|                                           |                      |
+-----+-----+
+-----+
| Processes: |
| GPU  GI   CI        PID   Type   Process name          GPU Memory |
|          ID   ID                                 Usage          |
+-----+-----+
| No running processes found |
+-----+-----+
+-----+

```

Step 10 From the Graid Technology website, download the latest version of the installer and make it executable.

Drivers & Documentation

```
$ sudo chmod +x <filename>
```

Driver Packages

| Model | x86_64 |
|---------|---|
| SR-1000 | graid-sr-installer-1.5.0-XXXXXXXXXX.run |
| SR-1001 | graid-sr-installer-1.5.0-XXXXXXXXXX.run |
| SR-1010 | graid-sr-installer-1.5.0-XXXXXXXXXX.run |

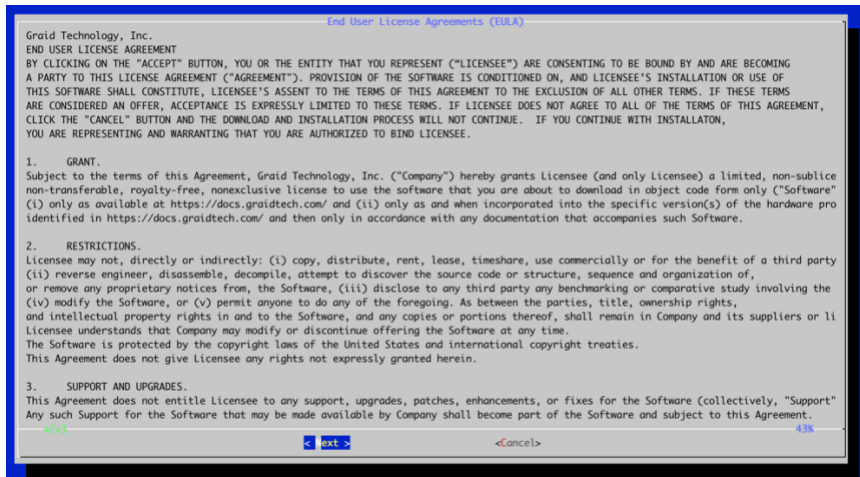
Step 11 Proceed to [Executing the Installer and Completing the Installation](#) on page 41 to execute the installer and to complete the installation.

Executing the Installer and Completing the Installation

Step 1 At the Welcome page select **Next** and click **Enter** to view the end-user license agreement.




Step 2 In the end-user license agreement, use the spacebar to scroll through the content. When you complete your review, select **Next** and click **Enter** to proceed.



Step 3 Type **accept**, click tab, select **Next**, and click **Enter** to accept the license agreement.



Step 4 Complete the installation, and the installer will reboot the system.



Step 5 To activate the software, apply the SupremeRAID™ license key.

```
$ sudo graidctl apply license <LICENSE_KEY>
```

USING THE SUPREME RAID™ DRIVER

This section describes how to use the basic functions of SupremeRAID™. It consists of step-by-step examples and command instructions that guide you to accessing all SupremeRAID™ features.

- To activate the SupremeRAID™ service, see [Activating the SupremeRAID™ Driver and Managing the License\(s\)](#) on page 43.
- To set up a local volume(Virtual Drive), see [Creating a RAID-5 Virtual Drive with Five NVMe SSDs](#) on page 45
- To set up an Initiator server, see [Creating a Physical Drive from the Remote NVMe-oF Targets](#) on page 48.
- To set up a Target server, see [Exporting the Virtual Drive as an NVMe-oF Target Drive Using RDMA to the Initiator](#) on page 47-49.
- To set up the high availability (HA) feature in one server, see [Setting Up the Dual-Controller to Enable HA and Auto-Failover](#) on page 48.

Activating the SupremeRAID™ Driver and Managing the License(s)

When you install the SupremeRAID™ driver, you must activate the SupremeRAID™ service by applying a specific license key prior to use the SupremeRAID™ service, and the license key you could get from your vendor. Once this is done, you can perform activities such as creating drive groups and virtual drives to use the SupremeRAID™.

- To check the SupremeRAID™ driver version, issue:

```
$ sudo graidctl version
```

- To activate the SupremeRAID™ software, issue:

```
$ sudo graidctl apply license <LICENSE_KEY>
```

- To check the license information, issue:

```
$ sudo graidctl describe license
```

- To check the controller status, issue:

```
$ sudo graidctl list controller
```

- To replace a new controller with the same model of the controller when the old controller is failure or missing, issue:

```
$ sudo graidctl disable controller <Controller_ID>
```

```
$ sudo graidctl replace controller <Controller_ID> <LICENSE_KEY>
```

- To delete the old controller that failed, missing, or disabled, issue:

```
$ sudo graidctl delete controller <Controller_ID>
```

Example

```
root@graid:/home/graid# graidctl version
✓Graidctl version successfully.
graidctl version:      1.5.0-rc1-644.ga82e0d9a.010
graid_server version:  1.5.0-rc1-644.ga82e0d9a.010
```

```
[graid@graid demo~]$ sudo graidctl apply license XXXXXXXX-XXXXXXXX-XXXXXXXX-XXXXXXXX
✓Apply license successfully.
[graid@graid demo~]$ sudo graidctl describe license
✓Describe license successfully.
License State:      APPLIED
Controller 0:
  Name: SR-1000
  Serial Number: 1xxxxxxxxx0
  License State: APPLIED
  License Key: XXXXXXXX-XXXXXXXX-XXXXXXXX-XXXXXXXX
  License Type: Full
  Expiration Days: Unlimited
  NVMe / NVMe-oF PD Number: 32
Controller 1:
  Name: SR-1000
  Serial Number: 1xxxxxxxxx1
  License State: APPLIED
  License Key: XXXXXXXX-XXXXXXXX-XXXXXXXX-XXXXXXXX
  License Type: Full
  Expiration Days: Unlimited
  NVMe / NVMe-oF PD Number: 32
Features:
  NVMe / NVMe-oF PD Number: 32
  RAID5: true
  RAID6: true
  Export VD via NVMe-oF: true
  Multiple Controller Support: true
[graid@graid demo~]$ sudo graidctl list controller
✓List controller successfully.


| ID | CONTROLLER MODEL | SERIAL NUMBER | NUMA | STATE  | DG  |
|----|------------------|---------------|------|--------|-----|
| 0  | SR-1000          | 1xxxxxxxxx0   | 0    | ONLINE | 0,1 |
| 1  | SR-1000          | 1xxxxxxxxx1   | 1    | ONLINE | 2,3 |


[graid@graid demo~]$ sudo graidctl disable controller 1
✓Disable controller successfully.
✓Disable controller Controller 1 successfully.
[graid@graid demo~]$ sudo graidctl disable controller 1
✓Disable controller successfully.
✓Disable controller Controller 1 successfully.
[graid@graid demo~]$ sudo graidctl replace controller 1 XXXXXXXX-XXXXXXXX-XXXXXXXX-XXXXXXXX
✓Replace controller successfully.
✓Replace controller Controller 1 successfully.
```

Note: When applying the license, you might need to provide the serial number of the NVIDIA GPU to Graid Technology Technical Support. To obtain the NVIDIA GPU serial number, issue either of the following commands. These commands list all NVIDIA cards in your environment and their serial number:

```
$ sudo nvidia-smi --query-gpu=name,index,serial --format=csv
```

OR

```
$ sudo nvidia-smi -q | grep -i serial
```

Note: If two controllers are activated in the `graid.conf` system configuration file, the SupremeRAID™ service prevents you from activating any additional controllers until one of the existing controllers is removed. This safeguard prevents conflicts and ensures proper system operation. Exercise caution and consult the software documentation or seek professional assistance if needed.

Creating a RAID-5 Virtual Drive with Five NVMe SSDs

To create a RAID-5 virtual drive with 5 NVMe SSDs:

Step 1 Create a physical drive.

```
$ sudo graidctl create physical_drive /dev/nvme0-4
```

Step 2 Create a drive group.

```
$ sudo graidctl create drive_group raid5 0-4
```

Step 3 Create a virtual drive with a 5TB volume size.

```
$ sudo graidctl create virtual_drive 0 5T
```

Step 4 Check the device path of the new virtual drive.

```
$ sudo graidctl list virtual_drive --dg-id=0
```

Output Example

```
[graid@graid demo~]$ sudo graidctl create physical_drive /dev/nvme0-4
✓Create physical drive successfully.
✓Create physical drive PD0 (/dev/nvme0: nqn.2019-08.org.qemu:NVME0002) successfully.
✓Create physical drive PD1 (/dev/nvme1: nqn.2019-08.org.qemu:NVME0004) successfully.
✓Create physical drive PD2 (/dev/nvme2: nqn.2019-08.org.qemu:NVME0001) successfully.
✓Create physical drive PD3 (/dev/nvme3: nqn.2019-08.org.qemu:NVME0003) successfully.
✓Create physical drive PD4 (/dev/nvme4: nqn.2019-08.org.qemu:NVME0005) successfully.
[graid@graid demo~]$ sudo graidctl create drive_group raid5 0-4
✓Create drive group successfully.
✓Create drive group DG0 successfully.
[graid@graid demo~]$ sudo graidctl create virtual_drive 0 5T
✓Create virtual drive successfully.
✓Create virtual drive DG0/VD0 successfully.
[graid@graid demo~]$ sudo graidctl list virtual_drive --dg-id=0
✓List virtual drive successfully.
```

| VD ID | DG ID | SIZE | DEVICE PATH | STATE | EXPORTED |
|-------|-------|---------|-------------|---------|----------|
| 0 | 0 | 4.7 TiB | /dev/gdg0n1 | OPTIMAL | No |

Creating a Physical Drive from the Remote NVMe-oF Targets

To create a physical drive from the Remote NVMe-oF targets:

Step 1 Connect to the remote NVMe-oF target.

```
$ sudo graidctl connect remote_target <tcp|rdma|fc> <addr> <address family> <service id>
```

Step 2 Check the NVMe drives from the remote NVMe-oF target.

```
$ sudo graidctl list nvme_drive
```

Step 3 Create the physical drives.

```
$ sudo graidctl create physical_drive <nqn or devpath>...
```

Step 4 Create a RAID5 drive group with four physical drives.

```
$ sudo graidctl create drive_group <Mode> <PD_ID>... [flags]
```

Output Example

```
[graid@graid demo~]$ sudo graidctl connect remote_target tcp 172.16.11.81 ipv4 4420
✓Connect remote target successfully.
✓Connect remote target Target 0 successfully.
[graid@graid demo~]$ sudo graidctl list nvme_drive
✓List nvme drive successfully.
```

| DEVICE PATH (4) | MODEL | NQN/wwid | NSID | CAPACITY | ADDRESS |
|-----------------|-------|---|------|----------|----------------------------------|
| /dev/nvme0n1 | Linux | uuid.b951d877-76af-4dfe-84ee-a45164554fe2 | 1 | 22 GB | traddr=172.16.11.81,trsvcid=4420 |
| /dev/nvme1n1 | Linux | uuid.6f21ec8f-00ee-4a30-a9b8-413447b8f138 | 1 | 22 GB | traddr=172.16.11.81,trsvcid=4420 |
| /dev/nvme2n1 | Linux | uuid.34d1d6aa-41fc-4c02-a660-f75429d7d74b | 1 | 22 GB | traddr=172.16.11.81,trsvcid=4420 |
| /dev/nvme3n1 | Linux | uuid.d846f451-31af-49ae-b3db-8ca90f454c3b | 1 | 22 GB | traddr=172.16.11.81,trsvcid=4420 |

```
[graid@graid demo~]$ sudo graidctl create physical_drive uuid.b951d877-76af-4dfe-84ee-a45164554fe2 /dev/nvme1 /dev/nvme3 uuid.34d1d6aa-41fc-4c02-a660-f75429d7d74b
✓Create physical drive successfully.
✓Create physical drive PD0 (uuid.b951d877-76af-4dfe-84ee-a45164554fe2) successfully.
✓Create physical drive PD1 (/dev/nvme1: uuid.6f21ec8f-00ee-4a30-a9b8-413447b8f138) successfully.
✓Create physical drive PD2 (/dev/nvme3: uuid.d846f451-31af-49ae-b3db-8ca90f454c3b) successfully.
✓Create physical drive PD3 (uuid.34d1d6aa-41fc-4c02-a660-f75429d7d74b) successfully.
[graid@graid demo~]$ sudo graidctl create drive_group raid5 0-3
✓ Create drive group DG0 successfully.
```


Exporting the Virtual Drive as an NVMe-oF Target Drive Using RDMA to the Initiator

To export the virtual drive as an NVMe-oF target drive using RDMA to the initiator:

Step 1 Create the RDMA/TCP NVMe-oF target port services.

```
$ sudo graidctl create nvmeof_target <tcp|rdma> <interface> <address family> <srvcid> [flags]
```

Step 2 Export a virtual drive as an NVMe-oF target.

```
$ sudo graidctl export virtual_drive <DG_ID> <VD_ID>... [flags]
```

Step 3 List all NVMe-oF targets.

```
$ sudo graidctl list nvmeof_target [flags]
```

Step 4 Describe the detailed information for an NVMe-oF target.

```
$ sudo graidctl describe nvmeof_target <PORT_ID> [flags]
```

Output Example

```
root@graid:/home/graid# graidctl ls virtual_drive
✓List virtual drive successfully.
```

| VD ID | DG ID | SIZE | DEVICE PATH | STATE | EXPORTED |
|-------|-------|---------|-------------|---------|----------|
| 0 | 0 | 959 MiB | /dev/gdg0n1 | OPTIMAL | No |

```
root@graid:/home/graid# graidctl create nvmeof_target tcp enp0s1 ipv4 4420
✓Create nvmeof target successfully.
✓Create nvmeof target Target 0 successfully.
root@graid:/home/graid# graidctl export virtual_drive 0 0 -i 0
✓Export virtual drive successfully.
✓Export virtual drive VD0 into Target 0 successfully.
root@graid:/home/graid# graidctl list nvmeof_target
✓List nvmeof target successfully.
```

| ID | TYPE | INTERFACE | ADDRESS | ADDRESS FAMILY | SERVICE ID | SUBSYSTEMS |
|----|------|-----------|-------------|----------------|------------|------------|
| 0 | tcp | enp0s1 | 172.16.55.5 | ipv4 | 4420 | DG0/VD0 |

```
root@graid:/home/graid# graidctl describe nvmeof_target 0
✓Describe nvmeof target successfully.
Id: 0
Port: 0
TransportType: tcp
Address: 172.16.55.5
Interface: enp0s1
AddressFamily: ipv4
ServiceId: 4420
Subsystems:
```

| NAME | DG ID | VD ID | DEVICE PATH |
|---|-------|-------|-------------|
| nqn.2020-05.com.graidtech:GRAID-SR4CCE211857D5F340:dg0vd0 | 0 | 0 | /dev/gdg0n1 |

Setting Up the Dual-Controller to Enable HA and Auto-Failover

To activate the HA feature, you need two SupremeRAID™ cards installed in your server model and have the service activated. The total drive group count is four, with at least one drive group allocated to each controller. However, the number of drive groups assigned to each controller does not need to be equal.

If one controller fails and the auto-failover function is turned on (it is enabled by default), the drive group under the failed controller fails over immediately to the functioning controller. To ensure data integrity, the drive group statuses that failover switch to Resync mode.

Step 1 Activate two cards to enable the HA feature.

```
$ sudo graidctl apply license <LICENSE_KEY>
```

Step 2 Check the controller status.

```
$ sudo graidctl list controller
```

Step 3 Check the NVMe devices' NUMA location.

```
$ sudo graidctl list nvme_drive -n <NUMA_ID>
```

Step 4 Create physical drives.

```
$ sudo graidctl create physical_drive <DEVICE_PATH|NQN|WWID>
```

Step 5 Create two drive groups with specific controllers.

```
$ sudo graidctl create drive_group <RAID_MODE> <PD_IDs> -c <Controller_ID>
```

Step 6 Create a specific virtual drive with a different drive group.

```
$ sudo graidctl create virtual_drive <DG_ID> [<VD_SIZE>]
```

Step 7 The drive group can optionally be assigned to a specific controller by editing it.

```
$ sudo graidctl edit <DG_ID> controller <Controller_ID>
```

Note: Typically, there is no need to set the controller manually while creating a drive group because SupremeRAID™ selects the optimal controller automatically based on the chosen physical drive. However, it is possible to adjust the controller manually for the drive group by making edits to it.

Output Example

```
[graid@graid demo~]$ sudo graidctl apply license XXXXXXXX-XXXXXXXX-XXXXXXXX-XXXXXXXX
✓Apply license successfully.
[graid@graid demo~]$ sudo graidctl apply license YYYYYYYY-YYYYYYYY-YYYYYYYY-XXXXXXXX
✗Apply license failed: New license PD number 12 is less than old license PD number 32
[graid@graid demo~]$ sudo graidctl apply license YYYYYYYY-YYYYYYYY-YYYYYYYY-YYYYYYYY
✓Apply license successfully.
[graid@graid demo~]$ sudo graidctl list controller
✓List controller successfully.
```

| ID | CONTROLLER MODEL | SERIAL NUMBER | NUMA | STATE | DG |
|----|------------------|---------------|------|--------|----|
| 0 | SR-1000 | 1xxxxxxxxxxx1 | 0 | ONLINE | |
| 1 | SR-1000 | 1xxxxxxxxxxx2 | 1 | ONLINE | |

```
[graid@graid demo~]$ sudo graidctl list nvme_drive -n 0
✓List nvme drive successfully.
```

| DEVICE PATH (3) | MODEL | NQN/NWID | NSID | CAPACITY | NUMA NODE | ADDRESS |
|-----------------|--------------|--|------|----------|-----------|--------------|
| /dev/nvme0n1 | KCM61VUL3T20 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A04HT1L8 | 1 | 50 GiB | 0 | 0000:22:00.0 |
| /dev/nvme2n1 | KCM61VUL3T20 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z0F0A031T1L8 | 1 | 50 GiB | 0 | 0000:23:00.0 |
| /dev/nvme4n1 | KCM61VUL3T20 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A058T1L8 | 1 | 50 GiB | 0 | 0000:41:00.0 |

```
[graid@graid demo~]$ sudo graidctl list nvme_drive -n 1
✓List nvme drive successfully.
```

| DEVICE PATH (3) | MODEL | NQN/NWID | NSID | CAPACITY | NUMA NODE | ADDRESS |
|-----------------|--------------|--|------|----------|-----------|--------------|
| /dev/nvme1n1 | KCM61VUL3T20 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z060A006T1L8 | 1 | 50 GiB | 1 | 0000:22:00.0 |
| /dev/nvme3n1 | KCM61VUL3T20 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A058T1L8 | 1 | 50 GiB | 1 | 0000:23:00.0 |
| /dev/nvme5n1 | KCM61VUL3T20 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z010A002T1L8 | 1 | 50 GiB | 1 | 0000:41:00.0 |

```
[graid@graid demo~]$ sudo graidctl create physical_drive /dev/nvme0,2,4
✓Create physical drive successfully.
✓Create physical drive PD0 (/dev/nvme0: nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A04HT1L8) successfully.
✓Create physical drive PD1 (/dev/nvme2: nqn.2019-10.com.kioxia:KCM61VUL3T20:Z0F0A031T1L8) successfully.
✓Create physical drive PD2 (/dev/nvme4: nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A058T1L8) successfully.
[graid@graid demo~]$ sudo graidctl create physical_drive /dev/nvme1,3,5
✓Create physical drive successfully.
✓Create physical drive PD3 (/dev/nvme1: nqn.2019-10.com.kioxia:KCM61VUL3T20:Z060A006T1L8) successfully.
✓Create physical drive PD4 (/dev/nvme3: nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A058T1L8) successfully.
✓Create physical drive PD5 (/dev/nvme5: nqn.2019-10.com.kioxia:KCM61VUL3T20:Z010A002T1L8) successfully.
[graid@graid demo~]$ sudo graidctl list physical_drive
✓List physical drive successfully.
```

| PD ID (6) | DG ID | DEVICE PATH | NQN/NWID | MODEL | CAPACITY | SLOT ID | NUMA NODE | STATE |
|-----------|-------|-------------|--|--------------|----------|---------|-----------|-------------------|
| 0 | N/A | /dev/gpd0 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A04HT1L8 | KCM61VUL3T20 | 50 GiB | N/A | 0 | UNCONFIGURED_GOOD |
| 1 | N/A | /dev/gpd1 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z0F0A031T1L8 | KCM61VUL3T20 | 50 GiB | N/A | 0 | UNCONFIGURED_GOOD |
| 2 | N/A | /dev/gpd2 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A058T1L8 | KCM61VUL3T20 | 50 GiB | N/A | 0 | UNCONFIGURED_GOOD |
| 3 | N/A | /dev/gpd3 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z060A006T1L8 | KCM61VUL3T20 | 50 GiB | N/A | 1 | UNCONFIGURED_GOOD |
| 4 | N/A | /dev/gpd4 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A058T1L8 | KCM61VUL3T20 | 50 GiB | N/A | 1 | UNCONFIGURED_GOOD |
| 5 | N/A | /dev/gpd5 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z010A002T1L8 | KCM61VUL3T20 | 50 GiB | N/A | 1 | UNCONFIGURED_GOOD |

```
[graid@graid demo~]$ sudo graidctl create drive_group raid5 0-2 -c 0
✓Create drive group successfully.
✓Create drive group DG0 successfully.
[graid@graid demo~]$ sudo graidctl list drive_group
✓List drive group successfully.
```

| DG ID | MODE | VD NUM | CAPACITY | FREE | USED | CONTROLLER | STATE |
|-------|-------|--------|----------|---------|------|----------------------|---------|
| 0 | RAID5 | 0 | 100 GiB | 100 GiB | 0 B | running: 0 prefer: 0 | OPTIMAL |

```
[graid@graid demo~]$ sudo graidctl create drive_group raid5 3-5 -c 1
✓Create drive group successfully.
✓Create drive group DG1 successfully.
[graid@graid demo~]$ sudo graidctl list drive_group
✓List drive group successfully.
```

| DG ID | MODE | VD NUM | CAPACITY | FREE | USED | CONTROLLER | STATE |
|-------|-------|--------|----------|---------|------|----------------------|---------|
| 0 | RAID5 | 0 | 100 GiB | 100 GiB | 0 B | running: 0 prefer: 0 | OPTIMAL |
| 1 | RAID5 | 0 | 100 GiB | 100 GiB | 0 B | running: 1 prefer: 1 | OPTIMAL |

```
[graid@graid demo~]$ sudo graidctl edit drive_group 1 controller 0
✓Edit drive group successfully.
[graid@graid demo~]$ sudo graidctl list drive_group
✓List drive group successfully.
```

| DG ID | MODE | VD NUM | CAPACITY | FREE | USED | CONTROLLER | STATE |
|-------|-------|--------|----------|---------|------|----------------------|---------|
| 0 | RAID5 | 0 | 100 GiB | 100 GiB | 0 B | running: 0 prefer: 0 | OPTIMAL |
| 1 | RAID5 | 0 | 100 GiB | 100 GiB | 0 B | running: 0 prefer: 0 | OPTIMAL |

```
[graid@graid demo~]$ sudo graidctl edit drive_group 1 controller 1
✓Edit drive group successfully.
[graid@graid demo~]$ sudo graidctl create virtual_drive 0
```

Upgrading the Software

To upgrade the Linux Driver, we offer two methods: silent upgrade and manual setup. Please follow the steps below for your preferred method.

Note: Perform the following procedure exactly as described. If you encounter any abnormal failure messages during the driver upgrade, please **collect the logs** and contact our Graid Technical Support team.

Silent Upgrade

In the SupremeRAID Linux Driver, if you have already installed the SupremeRAID driver, there's no need to uninstall it. Simply run the pre-installer and installer then include '—accept-license' in the upgrade command to automatically apply the license key to the new software.

Step 1 Download the upgrade driver package and make it executable.

Step 2 Run the pre-installer directly, it will automatically check the required dependencies.

```
$ sudo ./<filename>
```

Step 3 Run the installer and add 'accept-license' to automatically apply the license key.

```
$ sudo ./<filename> --accept-license
```

Step 4 Check the driver version to ensure the upgrade is successful.

```
$ sudo graidctl version
```

Manual Upgrade

If you need to perform a manual upgrade, please follow the steps below to upgrade the software.

Step 1 Stop all applications running on the virtual drive.

Step 2 Stop the management service.

```
$ sudo systemctl stop graid
```

Step 3 Make sure the SupremeRAID™ kernel module is unloaded.

```
$ sudo rmmod graid_nvidia graid
```

Step 4 Check the NVIDIA driver DKMS status.

```
$ sudo dkms status nvidia
```

Step 5 The version of the NVIDIA driver installed in the kernel must match the SupremeRAID™ driver version. If they do not match, perform the following steps to uninstall the NVIDIA driver.

A Dracut the initramfs (Centos, Rocky Linux, AlmaLinux, and RHEL).

```
$ sudo dracut --omit-drivers "nvidia graid" -f
```

B Uninstall the NVIDIA driver.

```
$ sudo ./usr/bin/nvidia-uninstall
```

C Install the new NVIDIA driver.

- Reboot the server.

Step 6 Uninstall the package using the command appropriate for your operating system.

- For Centos, Rocky Linux, AlmaLinux, RHEL, openSUSE, and SLES:

```
$ sudo rpm -e graid-sr
```

- For Ubuntu:

```
$ sudo dpkg -r graid-sr
```

Step 7 Confirm that the SupremeRAID™ module is unloaded. There should not be any output.

```
$ sudo lsmod | grep graid
```

Step 8 Confirm that the SupremeRAID™ package is uninstalled using the command appropriate for your operating system.

- For Centos, Rocky Linux, AlmaLinux, RHEL, openSUSE, and SLES:

```
$ sudo rpm -qa | grep graid
```

- For Ubuntu:

```
$ sudo dpkg -l | grep graid
```




There should not be any output.

Step 9 From the Graid Technology website, download the latest version of the installer and make it executable.

Drivers & Documentation

```
sudo chmod +x <filename>
```

Driver Packages

| Model | x86_64 |
|---------|--|
| SR-1000 | graid-sr-installer-1.5.0-.run |
| SR-1001 | graid-sr-installer-1.5.0-.run |
| SR-1010 | graid-sr-installer-1.5.0-.run |

Step 10 Start the graid service.

```
$ sudo systemctl enable graid
```

```
$ sudo systemctl start graid
```

Note: If you upgrade from version 1.2.x to version 1.5.x of the graid driver, the device path changes from /dev/gvdXn1 to /dev/gdgXnY.

Replacing a SupremeRAID™ Card

Step 1 Stop all applications running on the virtual drive.

Step 2 Stop the management service.

```
$ sudo systemctl stop graid
```

Step 3 Back up the configuration file.

```
$ sudo cp /etc/graid.conf graid.conf.bak
```

Step 4 Make sure the SupremeRAID™ kernel module is unloaded.

```
$ sudo rmmod graid_nvidia graid
```

Step 5 Check the NVIDIA driver DKMS status.

```
$ sudo dkms status nvidia
```

Note: The NVIDIA driver version installed in the kernel must match the graid driver version. Perform step 5 if the versions do not match.

Step 6 Uninstall the package using the command appropriate for your operating system:

- For Centos, Rocky Linux, AlmaLinux, RHEL, openSUSE, and SLES:

```
$ sudo rpm -e graid-sr
```

- For Ubuntu:

```
$ sudo dpkg -r graid-sr
```

Step 7 Confirm that the SupremeRAID™ module is unloaded.

```
$ sudo lsmod | grep graid
```

There should not be any output.

Step 8 Confirm that the SupremeRAID™ package is uninstalled using the command appropriate for your operating system.

- Centos, Rocky Linux, AlmaLinux, RHEL, openSUSE, and SLES:

```
$ sudo rpm -qa | grep graid
```

- Ubuntu:

```
sudo dpkg -l | grep graid
```

There should not be any output.

Step 9 Power-off the server, and then install the new card into the server.

Step 10 Power-on the server.

Step 11 From the Graid Technology website, download the latest version of the installer and make it executable.

Drivers & Documentation

```
$ sudo chmod +x <filename>
```

Step 12 When the installer finishes, restart the graidservice.

```
$ sudo systemctl restart graid
```

If the settings do not return properly after restarting graidservice, see [Manually Migrating the RAID Configuration Between Hosts](#) on page 127.

Note: If you are replacing a card in the system, deleting any inactive or invalid licenses associated with the old card is essential. Failing to do so may prevent other cards from becoming active, which is key for multi-controller systems.

COMMANDS AND SHORTCUTS

Syntax

Use the following syntax to run `graidctl` commands from the terminal window:

```
graidctl [command] [OBJECT_TYPE] [OBJECT_ID] [flags]
```

where `command`, `OBJECT_TYPE`, `OBJECT_ID`, and `flags` are:

- **command**: Specifies the operation to perform on one or more resources (for example create, list, describe, and delete).
- **OBJECT_TYPE**: Specifies the object type. Object types are case-sensitive (for example `license`, `physical_drive`, and `drive_group`).
- **OBJECT_ID**: Specifies the object ID. Some commands support simultaneous operations on multiple objects. You can specify the `OBJECT_ID` individually or use a dash to describe an `OBJECT_ID` range. For example, to delete physical drives 1, 3, 4, and 5 simultaneously, issue the command:

```
$ sudo graidctl delete physical_drive 1 3-5
```

- **flags**: Specifies optional flags. For example:

- **-force** forces the deletion of a physical drive.

```
$ sudo graidctl delete physical_drive 0 -force
```

- **-json** prints output in JSON format. This flag can also assist with API implementation.

```
$ sudo graidctl list virtual_drive --format json
```

For help, run `graidctl help` from the terminal window.

Command and Subcommand Quick Reference

General

| Category | Commands | Alias | Sub-Commands | alias |
|----------|----------|-------|--------------|-------|
| Common | version | | | |
| License | apply | | license | lic |
| | describe | desc | license | lic |

Resources

| Category | Commands | Alias | Sub-Commands | alias |
|----------------|----------|----------------|----------------|-------|
| NVMe Drive | list | l, ls | nvme_drive | nd |
| SCSi Drive | list | l, ls | scsi_drive | sd |
| Physical Drive | create | c, cre, crt | physical_drive | pd |
| | icreate | ic, icre, icrt | physical_drive | pd |
| | delete | d, del | physical_drive | pd |
| | describe | desc | physical_drive | pd |
| | edit | e | physical_drive | pd |
| | list | l, ls | physical_drive | pd |
| | replace | en | physical_drive | pd |
| Drive Group | create | c, cre, crt | drive_group | dg |
| | icreate | ic, icre, icrt | drive_group | dg |
| | delete | d, del | drive_group | dg |

| Category | Commands | Alias | Sub-Commands | alias |
|---------------|----------|----------------|---------------|-------|
| | describe | desc | drive_group | dg |
| | edit | e | drive_group | dg |
| | list | l, ls | drive_group | dg |
| Virtual Drive | create | c, cre, crt | virtual_drive | vd |
| | icreate | ic, icre, icrt | virtual_drive | vd |
| | delete | d, del | virtual_drive | vd |
| | describe | desc | virtual_drive | vd |
| | edit | e | virtual_drive | vd |
| | list | l, ls | virtual_drive | vd |
| Controller | enable | | controller | cx |
| | disable | | controller | cx |
| | delete | d, del | controller | cx |
| | list | l, ls | controller | cx |
| | replace | en | controller | cx |
| MD Boot Drive | import | im, imp | md_drive | md |
| | replace | en | md_drive | md |
| Config | describe | desc | config | conf |
| | edit | e | config | conf |
| | delete | d, del | config | conf |
| | restore | Re | Config | conf |
| Event | delete | d, del | event | ev |
| | list | l, ls | event | ev |

Features

| Category | Commands | Alias | Sub-Commands | alias |
|-------------------|------------|--------------|-------------------|-------|
| Consistency Check | describe | desc | consistency_check | cc |
| | set | | consistency_check | cc |
| | start | | consistency_check | cc |
| | stop | | consistency_check | cc |
| Export NVMe-oF | create | c, cre, crt | nvmeof_target | nt |
| | describe | desc | nvmeof_target | nt |
| | delete | d, del | nvmeof_target | nt |
| | list | l, ls | nvmeof_target | nt |
| | export | ex, exp | virtual_drive | vd |
| | unexport | unex, unexp | virtual_drive | vd |
| Import NVMe-oF | connect | conn | remote_target | rt |
| | disconnect | dis, disconn | remote_target | rt |
| | list | l, ls | remote_target | rt |

Managing Licenses

You can apply the license and check license information.

Applying the License

To apply the license and complete the installation, issue the following command:

```
$ sudo graidctl apply license <LICENSE_KEY> [flags]
```

OR

```
$ sudo graidctl apply lic <LICENSE_KEY> [flags]
```

| Flag | Description |
|------------|------------------------------------|
| -h, --help | Help for the apply license command |

Output example for invalid and valid licenses is shown below:

```
[graid@graid-demo ~]$ sudo graidctl apply license XXXXXXXX-XXXXXXXX-XXXXXXXX-XXXXXX0X
✖Apply license failed: Incorrect license format: XXXXXXXX-XXXXXXXX-XXXXXXXX-XXXXXX0X
[graid@graid-demo ~]$ sudo graidctl apply license XXXXXXXX-XXXXXXXX-XXXXXXXX-XXXXXXXX
✔Apply license successfully.
[graid@graid-demo ~]$ sudo graidctl apply lic XXXXXXXX-XXXXXXXX-XXXXXXXX-XXXXXXXX
✔Apply license successfully.
```

Note: When applying the license, you must provide the serial number of the NVIDIA GPU to Graid Technology Technical Support.

To obtain NVIDIA GPU serial number, issue the following command:

```
$ sudo nvidia-smi --query-gpu=name,index,serial --format=csv
```

OR

```
$ sudo nvidia-smi -q | grep -i serial
```

This command lists all NVIDIA cards in your environment and their serial number.

Checking License Information

To obtain the license information, issue the following command:

```
$ sudo graidctl describe license [flags]
```

OR

```
$ sudo graidctl desc lic [flags]
```

| Flag | Description |
|------------|---------------------------------------|
| -h, --help | Help for the describe license command |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl describe license
✓Describe license successfully.
Controller 0:
    Name: SR-1000
    Serial Number: 1352424094196
    License State: APPLIED
    License Key: XXXXXXXX-XXXXXXX-XXXXXXX-XXXXXXX
    License Type: Full
    Expiration Days: Unlimited
    NVMe / NVMe-oF PD Number: 32

Controller 1:
    Name: SR-1000
    Serial Number: 1320439569794
    License State: APPLIED
    License Key: XXXXXXXX-XXXXXXX-XXXXXXX-XXXXXXX
    License Type: Full
    Expiration Days: Unlimited
    NVMe / NVMe-oF PD Number: 32

Features:
    NVMe / NVMe-oF PD Number: 32
    RAID5: true
    RAID6: true
    Export VD via NVMe-oF: true
    Multiple Controller Support: true
```

Output Content

| Field | Description |
|--------------------------|---|
| Name | Product SKU |
| Serial Number | Applied controller's serial number |
| License State | License state (see the following table) |
| License Key | Applied license key |
| License Type | License type (Full or Essential) |
| Expiration Days | Expiration date of the license key |
| NVMe / NVMe-oF PD Number | This license allows for a maximum number of PDs for NVMe/NVMe-oF. |

License State

| State | Description |
|-----------|--|
| UNAPPLIED | License was not applied. |
| APPLIED | A valid license was applied. |
| INVALID | A valid license was applied, but a valid RAID card cannot be detected. |

Feature Support

| Features | Description | Value |
|-----------------------------|--|---------|
| NVMe / NVMe-oF PD Number | Accept total create maximum amount of the PD | Integer |
| RAID5 | Support RAID5 function | Boolean |
| RAID6 | Support RAID6 function | Boolean |
| Export VD via NVMe-oF | Support Export NVMe-of function | Boolean |
| Multiple Controller Support | Support Multiple Controller function | Boolean |

Checking the SupremeRAID™ Driver Version

You can prompt the version command to check graidservice information.

To obtain the graidservice version information, issue the following command:

```
$ sudo graidctl version [flags]
```

| Flag | Description |
|------------|---|
| -h, --help | Help for the check graidservice version command |

Output Example

```
root@graid:/home/graid# graidctl version
✓Graidctl version successfully.
graidctl version:      1.5.0-rc1-644.ga82e0d9a.010
graid_server version: 1.5.0-rc1-644.ga82e0d9a.010
```

Viewing Host Drive Information

Listing NVMe Drives

To list all the directly attached NVMe drives or NVMe-oF target drives that can be used to create physical drives, issue the following command:

```
$ sudo graidctl list nvme_drive [flags]
```

OR

```
$ sudo graidctl ls nd [flags]
```

| Flag | Description |
|-----------------|--|
| -h, --help | Help for the list nvme_drive command |
| -n, --numa-node | [int32] Filter by numa node Default: -1 |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl list nvme_drive
✔List nvme drive successfully.
```

| DEVICE PATH(4) | MODEL | NQN/wwID | NSID | CAPACITY | NUMA NODE | ADDRESS |
|----------------|--------------|--|------|----------|-----------|--------------|
| /dev/nvme0 | KCM61VUL3T20 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A064T1L8 | 1 | 3.2 TB | 1 | 0000:e4:00.0 |
| /dev/nvme1 | KCM61VUL3T20 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z050A002T1L8 | 1 | 3.2 TB | 0 | 0000:01:00.0 |
| /dev/nvme2 | KCM61VUL3T20 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A05KT1L8 | 1 | 3.2 TB | 1 | 0000:e1:00.0 |
| /dev/nvme3 | KCM61VUL3T20 | nqn.2019-10.com.kioxia:KCM61VUL3T20:X0N0A015T1L8 | 1 | 3.2 TB | 0 | 0000:43:00.0 |

```
[graid@graid-demo ~]$ sudo graidctl ls nd
✔List nvme drive successfully.
```

| DEVICE PATH(4) | MODEL | NQN/wwID | NSID | CAPACITY | NUMA NODE | ADDRESS |
|----------------|--------------|--|------|----------|-----------|--------------|
| /dev/nvme0 | KCM61VUL3T20 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A064T1L8 | 1 | 3.2 TB | 1 | 0000:e4:00.0 |
| /dev/nvme1 | KCM61VUL3T20 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z050A002T1L8 | 1 | 3.2 TB | 0 | 0000:01:00.0 |
| /dev/nvme2 | KCM61VUL3T20 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A05KT1L8 | 1 | 3.2 TB | 1 | 0000:e1:00.0 |
| /dev/nvme3 | KCM61VUL3T20 | nqn.2019-10.com.kioxia:KCM61VUL3T20:X0N0A015T1L8 | 1 | 3.2 TB | 0 | 0000:43:00.0 |

```
[graid@graid-demo ~]$ sudo graidctl ls nd -n 1
✔List nvme drive successfully.
```

| DEVICE PATH(2) | MODEL | NQN/wwID | NSID | CAPACITY | NUMA NODE | ADDRESS |
|----------------|--------------|--|------|----------|-----------|--------------|
| /dev/nvme0 | KCM61VUL3T20 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A064T1L8 | 1 | 3.2 TB | 1 | 0000:e4:00.0 |
| /dev/nvme2 | KCM61VUL3T20 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A05KT1L8 | 1 | 3.2 TB | 1 | 0000:e1:00.0 |

Output Content

| Field | Description |
|-------------|----------------------------------|
| DEVICE PATH | Block device path of the drive |
| NQN | NVMe Qualified Name of the drive |
| MODEL | Model number of the drive |
| CAPACITY | Capacity of the drive |
| NUMA MODE | NUMA NODE of the drive |

Listing SAS/SATA Drives

To list all SAS/SATA drives that can be used as physical drives, issue the following command:

```
$ sudo graidctl list scsi_drive
```

OR

```
$ sudo graidctl ls sd
```

| Flag | Description |
|------------|--------------------------------------|
| -h, --help | Help for the list scsi_drive command |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl list scsi_drive
✔List scsi drive successfully.
```

| DEVICE PATH | WWID | MODEL | CAPACITY |
|-------------|--|------------------|----------|
| /dev/sda | t10.ATA INTEL SSDSC2KB240G7 BTYS83010GKS240AGN | INTEL SSDSC2KB24 | 240 GB |
| /dev/sdb | t10.ATA INTEL SSDSC2KB240G8 BTYF052107VH240AGN | INTEL SSDSC2KB24 | 240 GB |

```
[graid@graid-demo ~]$ sudo graidctl ls sd
✔List scsi drive successfully.
```

| DEVICE PATH | WWID | MODEL | CAPACITY |
|-------------|--|------------------|----------|
| /dev/sda | t10.ATA INTEL SSDSC2KB240G7 BTYS83010GKS240AGN | INTEL SSDSC2KB24 | 240 GB |
| /dev/sdb | t10.ATA INTEL SSDSC2KB240G8 BTYF052107VH240AGN | INTEL SSDSC2KB24 | 240 GB |

Output Content

| Field | Description |
|-------------|---------------------------------------|
| DEVICE PATH | Block device path of the drive |
| WWID | Worldwide Identification of the drive |
| MODEL | Model number of the drive |
| CAPACITY | Capacity of the drive |

Managing Physical Drives

Creating a Physical Drive

To create a physical drive, issue the following command:

```
$ sudo graidctl create physical_drive <DEVICE_PATH|NQN|WWID> [flag]
```

OR

```
sudo graidctl c pd <DEVICE_PATH|NQN|WWID> [flag]
```

| Flag | Description |
|--------------|--|
| -h, --help | Help for the list physical_drive command |
| -f, --dblfwd | Door Bell Forwarding |

Output Example

The following figure shows an output example when creating multiple physical drives simultaneously with the device path and NQN.

```
[graid@graid-demo ~]$ sudo graidctl create physical_drive /dev/nvme0-3
✓Create physical drive successfully.
✓Create physical drive PD0 (/dev/nvme0: nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A064T1L8) successfully.
✓Create physical drive PD1 (/dev/nvme1: nqn.2019-10.com.kioxia:KCM61VUL3T20:Z050A002T1L8) successfully.
✓Create physical drive PD2 (/dev/nvme2: nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A05KT1L8) successfully.
✓Create physical drive PD3 (/dev/nvme3: nqn.2019-10.com.kioxia:KCM61VUL3T20:X0N0A015T1L8) successfully.
[graid@graid-demo ~]$ sudo graidctl create physical_drive nqn.2019-10.com.kioxia:KCM61VUL3T20:X0X0A01ET1L8 \
> nqn.2019-10.com.kioxia:KCM61VUL3T20:Z0F0A031T1L8
✓Create physical drive PD8 (nqn.2019-10.com.kioxia:KCM61VUL3T20:X0X0A01ET1L8) successfully.
✓Create physical drive PD9 (nqn.2019-10.com.kioxia:KCM61VUL3T20:Z0F0A031T1L8) successfully.
[graid@graid-demo ~]$ sudo graidctl c pd /dev/nvme4,7,8
✓Create physical drive successfully.
✓Create physical drive PD10 (/dev/nvme4: nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A032T1L8) successfully.
✓Create physical drive PD11 (/dev/nvme7: nqn.2019-10.com.kioxia:KCM61VUL3T20:Z050A078T1L8) successfully.
✓Create physical drive PD12 (/dev/nvme8: nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A09XT1L8) successfully.
```

Note: Be sure the system or other applications are not on the physical drive before creating or replacing the drive.

Listing the Physical Drives

To list all of the physical drives, issue the following command:

```
$ sudo graidctl list physical_drive [flag]
```

OR

```
$ sudo graidctl ls pd [flag]
```

| Flag | Description |
|----------------|--|
| -h, --help | Help for the list physical_drive command |
| -d, --dg-id | [int32] Filter result by drive group ID Default: -1 |
| -f, --free | List unused PDs |
| -l, --locating | List locating PDs |

| Flag | Description |
|-----------------|--|
| -n, --numa-node | [int32] Filter by numa node Default: -1 |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl list physical drive
✔List physical drive successfully.
```

| PD ID (10) | DG ID | DEVICE PATH | MQN/WNID | MODEL | CAPACITY | SLOT ID | NUMA NODE | STATE |
|------------|-------|--------------|--|--------------|----------|---------|-----------|-------------------|
| 0 | N/A | /dev/gpd0 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A038T1L8 | KCM61VUL3T20 | 3.2 TB | 0 | 1 | UNCONFIGURED_GOOD |
| 1 | N/A | /dev/gpd1 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A060T1L8 | KCM61VUL3T20 | 3.2 TB | 1 | 0 | UNCONFIGURED_GOOD |
| 2 | N/A | /dev/gpd2 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A04WT1L8 | KCM61VUL3T20 | 3.2 TB | 2 | 1 | UNCONFIGURED_GOOD |
| 3 | N/A | /dev/gpd3 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z050A002T1L8 | KCM61VUL3T20 | 3.2 TB | 3 | 0 | UNCONFIGURED_GOOD |
| 4 | N/A | /dev/gpd4 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z010A003T1L8 | KCM61VUL3T20 | 3.2 TB | 4 | 1 | UNCONFIGURED_GOOD |
| 5 | N/A | /dev/gpd5 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z060A005T1L8 | KCM61VUL3T20 | 3.2 TB | 5 | 0 | UNCONFIGURED_GOOD |
| 6 | N/A | /dev/gpd6 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z0F0A031T1L8 | KCM61VUL3T20 | 3.2 TB | 6 | 1 | UNCONFIGURED_GOOD |
| 7 | N/A | /dev/gpd7 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z010A002T1L8 | KCM61VUL3T20 | 3.2 TB | 7 | 0 | UNCONFIGURED_GOOD |
| 32 | 4 | /dev/nvme0n1 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A04HT1L8 | KCM61VUL3T20 | 3.2 TB | N/A | 1 | ONLINE |
| 33 | 4 | /dev/nvme1n1 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z010A001T1L8 | KCM61VUL3T20 | 3.2 TB | N/A | 1 | ONLINE |

```
[graid@graid-demo ~]$ sudo graidctl ls pd
✔List physical drive successfully.
```

| PD ID (10) | DG ID | DEVICE PATH | MQN/WNID | MODEL | CAPACITY | SLOT ID | NUMA NODE | STATE |
|------------|-------|--------------|--|--------------|----------|---------|-----------|-------------------|
| 0 | N/A | /dev/gpd0 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A038T1L8 | KCM61VUL3T20 | 3.2 TB | 0 | 1 | UNCONFIGURED_GOOD |
| 1 | N/A | /dev/gpd1 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A060T1L8 | KCM61VUL3T20 | 3.2 TB | 1 | 0 | UNCONFIGURED_GOOD |
| 2 | N/A | /dev/gpd2 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A04WT1L8 | KCM61VUL3T20 | 3.2 TB | 2 | 1 | UNCONFIGURED_GOOD |
| 3 | N/A | /dev/gpd3 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z050A002T1L8 | KCM61VUL3T20 | 3.2 TB | 3 | 0 | UNCONFIGURED_GOOD |
| 4 | N/A | /dev/gpd4 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z010A003T1L8 | KCM61VUL3T20 | 3.2 TB | 4 | 1 | UNCONFIGURED_GOOD |
| 5 | N/A | /dev/gpd5 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z060A005T1L8 | KCM61VUL3T20 | 3.2 TB | 5 | 0 | UNCONFIGURED_GOOD |
| 6 | N/A | /dev/gpd6 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z0F0A031T1L8 | KCM61VUL3T20 | 3.2 TB | 6 | 1 | UNCONFIGURED_GOOD |
| 7 | N/A | /dev/gpd7 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z010A002T1L8 | KCM61VUL3T20 | 3.2 TB | 7 | 0 | UNCONFIGURED_GOOD |
| 32 | 4 | /dev/nvme0n1 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A04HT1L8 | KCM61VUL3T20 | 3.2 TB | N/A | 1 | ONLINE |
| 33 | 4 | /dev/nvme1n1 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z010A001T1L8 | KCM61VUL3T20 | 3.2 TB | N/A | 1 | ONLINE |

```
[graid@graid-demo ~]$ sudo graidctl ls pd -n 0
✔List physical drive successfully.
```

| PD ID (4) | DG ID | DEVICE PATH | MQN/WNID | MODEL | CAPACITY | SLOT ID | NUMA NODE | STATE |
|-----------|-------|-------------|--|--------------|----------|---------|-----------|-------------------|
| 1 | N/A | /dev/gpd1 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A060T1L8 | KCM61VUL3T20 | 3.2 TB | 1 | 0 | UNCONFIGURED_GOOD |
| 3 | N/A | /dev/gpd3 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z050A002T1L8 | KCM61VUL3T20 | 3.2 TB | 3 | 0 | UNCONFIGURED_GOOD |
| 5 | N/A | /dev/gpd5 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z060A005T1L8 | KCM61VUL3T20 | 3.2 TB | 5 | 0 | UNCONFIGURED_GOOD |
| 7 | N/A | /dev/gpd7 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z010A002T1L8 | KCM61VUL3T20 | 3.2 TB | 7 | 0 | UNCONFIGURED_GOOD |

Output Content

| Field | Description |
|---------|---|
| SLOT ID | Slot ID of the corresponding NVMe/SAS/SATA drive. The PD ID is not related to the SLOT ID. To set the physical drives, use the PD ID. |
| DG ID | Drive group ID of the physical drive |

| Field | Description |
|----------|--|
| PD ID | PD ID. The PD ID is a unique ID provided by the SupremeRAID™ driver when the physical drive is created. It is not related to any SSD information such as slot ID or NQN. The PD ID is used for all further operations. |
| NQN/WWID | NQN or WWID of corresponding NVMe/SAS/SATA drive |
| MODEL | Model number of the corresponding NVMe/SAS/SATA drive |
| CAPACITY | Capacity of corresponding NVMe/SAS/SATA drive |
| NODE | NUMA NODE of the corresponding NVMe/SAS/SATA drive |
| STATE | State of the physical drive (see the following table). |

Physical Drive State

| State | Description |
|-------------------|---|
| ONLINE | Physical drive was added to a drive group and is ready to work. |
| HOTSPARE | Physical drive is configured as a hot spare drive. |
| FAILED | Physical drive is detected, but it is not operating normally. |
| OFFLINE | Physical drive is marked as offline. |
| REBUILD | Physical drive is being rebuilt. |
| MISSING | Physical drive cannot be detected. |
| INCONSISTENT | Data in the physical drive is inconsistent. This condition can occur when the physical drive is in the REBUILD state and the system encounters a crash. |
| UNCONFIGURED_GOOD | Physical drive did not join a drive group. |
| UNCONFIGURED_BAD | Physical drive did not join a drive group and is not operating normally. |

Deleting a Physical Drive

To delete a physical drive, issue the following command:

```
$ sudo graidctl delete physical_drive <PD_ID>
```

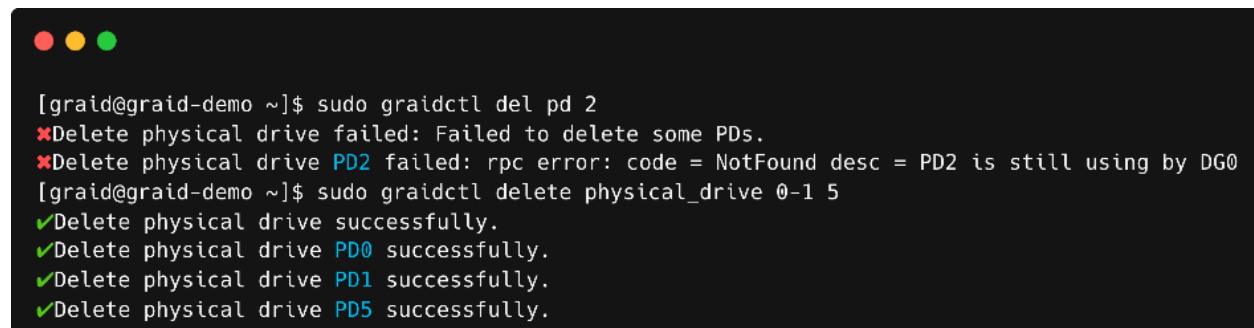
OR

```
$ sudo graidctl del pd <PD_ID>
```

| Flag | Description |
|------------|--|
| -h, --help | Help for the list physical_drive command |

Output Example

The following figure shows an output example for deleting multiple physical drives simultaneously.



```
[graid@graid-demo ~]$ sudo graidctl del pd 2
✗Delete physical drive failed: Failed to delete some PDs.
✗Delete physical drive PD2 failed: rpc error: code = NotFound desc = PD2 is still using by DG0
[graid@graid-demo ~]$ sudo graidctl delete physical_drive 0-1 5
✓Delete physical drive successfully.
✓Delete physical drive PD0 successfully.
✓Delete physical drive PD1 successfully.
✓Delete physical drive PD5 successfully.
```

The output shows that a physical drive cannot be deleted when it is part of a drive group.

Describing a Physical Drive

To view detailed information for a physical drive, issue the following command:

```
$ sudo graidctl describe physical_drive <PD_ID>
```

OR

```
$ sudo graidctl desc pd <PD_ID>
```


| Flag | Description |
|------------|--|
| -h, --help | Help for the describe physical_drive command |

Output Example:

```
[graid@graid-demo ~]$ sudo graidctl describe physical_drive 5
✓Describe physical drive successfully.
PD ID:          5
DG ID:          -1
Slot ID:        15
GUID:           nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A038T1L8
Mode:           KCM61VUL3T20
Capacity:       3.2 TB
State:          HOTSPARE
Device Path:    /dev/gpd8
Attributes:
                locating = false
                hotspare = 0,1
[graid@graid-demo ~]$ sudo graidctl desc pd 0
✓Describe physical drive successfully.
PD ID:          0
DG ID:          0
Slot ID:        1
GUID:           nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080IP38T1L8
Mode:           KCM61VUL3T20
Capacity:       3.2 TB
State:          HOTSPARE
Device Path:    /dev/gpd0
Attributes:
                locating = true
                hotspare =
```

Locating a Physical Drive

To locate a physical drive, issue the following command:

```
$ sudo graidctl edit physical_drive <PD_ID> locating start
```

To stop locating a physical drive, issue the following command:

```
$ sudo graidctl edit physical_drive <PD_ID> locating stop
```

Marking a Physical Drive Online or Offline

To mark a physical drive as online or offline, issue the following command:

```
$ sudo graidctl edit physical_drive <PD_ID> marker <offline|online>
```

Note: Marking a physical drive as offline, even briefly, puts the physical drive in the **REBUILD** state.

Assigning a Hot Spare Drive

To assign a physical drive as global hot spare, issue the following command:

```
$ sudo graidctl edit physical_drive <PD_ID> hotspare <DG_ID>
```

To assign a physical drive as the hot spare for a specific drive group, issue the following command:

```
$ sudo graidctl edit physical_drive <PD_ID> hotspare <DG_ID>
```

To assign a physical drive as a hot spare for multiple drive groups, use a comma (,) to separate the drive group IDs.

| Flag | Description |
|------------|--|
| -h, --help | Help for the edit physical_drive setting command |

Replacing a Nearly Worn-Out or Broken SSD

Note: Make sure the system or other applications are not on the physical drive before creating or replacing the drive.

To replace a nearly worn-out or broken SSD:

Step 1 If the physical drive is in the MISSING or other abnormal state, skip this step. Otherwise, issue the following command to mark the physical drive as bad:

```
$ sudo graidctl edit pd <OLD_PD_ID> marker bad
```

Step 2 Replace the NVMe SSD. The state of the prior physical drive indicates **FAILED**.

Step 3 Check the NQN of the new SSD.

```
$ sudo graidctl list nvme_drive
```

Step 4 Replace the physical drive.

```
$ sudo graidctl replace physical_drive <OLD_PD_ID> <DEVICE_PATH|NQN|WWID>
```

Output Example

```
[graid@graid demo ~]$ sudo graidctl edit physical_drive 0 marker bad
✓Edit physical drive successfully.
✓Edit physical drive PD0 successfully.
[graid@graid demo ~]$ sudo graidctl list physical_drive
✓List physical drive successfully.
```

| PD ID (5) | DG ID | DEVICE PATH | NQN/NWID | MODEL | CAPACITY | SLOT ID | STATE |
|-----------|-------|-------------|--|--------------|----------|---------|--------|
| 0 | 0 | /dev/gpd0 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z010A004T1L8 | KCM61VUL3T20 | 3.2 TB | 15 | FAILED |
| 1 | 0 | /dev/gpd1 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z060A006T1L8 | KCM61VUL3T20 | 3.2 TB | 9 | ONLINE |
| 2 | 0 | /dev/gpd2 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z010A001T1L8 | KCM61VUL3T20 | 3.2 TB | 8 | ONLINE |
| 3 | 0 | /dev/gpd3 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A04HT1L8 | KCM61VUL3T20 | 3.2 TB | 11 | ONLINE |
| 4 | 0 | /dev/gpd4 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A05KT1L8 | KCM61VUL3T20 | 3.2 TB | 3 | ONLINE |

```
[graid@graid demo ~]$ sudo graidctl list nvme_drive
✓List nvme drive successfully.
```

| DEVICE PATH (1) | NQN | MODEL | CAPACITY |
|-----------------|--|--------------|----------|
| /dev/nvme5 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z050A002T1L8 | KCM61VUL3T20 | 3.2 TB |

```
[graid@graid demo ~]$ sudo graidctl replace physical_drive 0 /dev/nvme5
✓Replace physical drive successfully.
[graid@graid demo ~]$ sudo graidctl list physical_drive
✓List physical drive successfully.
```

| PD ID (5) | DG ID | DEVICE PATH | NQN/NWID | MODEL | CAPACITY | SLOT ID | STATE |
|-----------|-------|-------------|--|--------------|----------|---------|-------------------------------------|
| 1 | 0 | /dev/gpd1 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z060A006T1L8 | KCM61VUL3T20 | 3.2 TB | 15 | ONLINE |
| 2 | 0 | /dev/gpd2 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z010A001T1L8 | KCM61VUL3T20 | 3.2 TB | 9 | ONLINE |
| 3 | 0 | /dev/gpd3 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A04HT1L8 | KCM61VUL3T20 | 3.2 TB | 8 | ONLINE |
| 4 | 0 | /dev/gpd4 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A05KT1L8 | KCM61VUL3T20 | 3.2 TB | 11 | ONLINE |
| 5 | 0 | /dev/gpd5 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z050A002T1L8 | KCM61VUL3T20 | 3.2 TB | 3 | REBUILD (12.69%, 54 mins remaining) |

```
[graid@graid demo ~]$ sudo graidctl list drive_group
✓List drive group successfully.
```

| DG ID | MODE | VD NUM | CAPACITY | FREE | USED | STATE |
|-------|-------|--------|----------|-------|--------|----------|
| 0 | RAID5 | 1 | 13 TB | 12 TB | 1.0 TB | RECOVERY |

```
[graid@graid demo ~]$ sudo graidctl list virtual_drive
✓List virtual drive successfully.
```

| VD ID | DG ID | SIZE | DEVICE PATH | STATE |
|-------|-------|--------|-------------|----------|
| 0 | 0 | 1.0 TB | /dev/gdg0n1 | RECOVERY |

Managing Drive Groups

Creating Drive Groups

To create a drive group or groups, issue the following command:

```
$ sudo graidctl create drive_group <RAID_MODE> <PD_IDs> [flag]
```

OR

```
$ sudo graidctl c dg <RAID_MODE> <PD_IDs> [flag]
```

| Flag | Description |
|-----------------------|--|
| -h, --help | Help for the create drive_group command |
| -b, --background-init | Background initialization |
| -c, --controller | [int32] Specific controller id Default: -1 |
| -f, --force-clean | Ignore initialization (Danger) |
| -z, --foreground-init | Foreground initialization (Write Zeros) |
| -s, --strip-size | [uint32] Strip Size (KiB) Values: 4, 8, 16, 32, 64, 128 Default: 4 |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl create drive_group raid1 0-1
✓Create drive group DG0 successfully.
[graid@graid-demo ~]$ sudo graidctl create drive_group raid5 2-4
✓Create drive group DG1 successfully.
[graid@graid-demo ~]$ sudo graidctl create drive_group raid6 5-9
✓Create drive group DG2 successfully.
```

Required Parameters

| Option | Description |
|-----------|---|
| RAID_MODE | RAID mode of the drive group. Entries must be all uppercase or all lowercase. For example, RAID6 or raid6 are both correct. |
| PD_IDS | ID of the physical drive joining the drive group. |

Optional Parameters

| Option | Description | Behavior |
|------------------------------|--|---|
| --background -init, -b | Default option. Use standard methods to initialize the drive group. When all the physical drives in the drive group support the de-allocate dataset management command, it is used to synchronize the data, or parity, between the physical drives during the creation of the drive group. | An I/O-capable device path similar to /dev/gdg0n1 is created. |
| --foreground -init, -z | Initializing foreground. This method writes zeros to whole drives | The virtual drive appears in the system after initialization is complete. Use the following command to check the initialization progress: <pre>\$ sudo graidctl list drive_group</pre> |
| --force -clean, -f | Force bypass initialize. Assumes that the drives are all clean. | The drive group STATE immediately becomes OPTIMAL, indicating that the drive group is available for use. |
| --controller, -c | Specific controller to control this drive_group. Default: -1, [Int32] | The drive group control by specific controller. |
| --strip-size, -s | Strip size of the drive_group. [RAID0,RAID10] Values: 4, 8, 16, 32, 64, 128 Default: 4, [Int32] | Adjust RAID0/RAID10 strip size to a specific size: (4k, 8k, 16k, 32k, 64k, or 128k) |

Wait for the drive group initialization to complete. DO NOT power-off or reboot the system when the drive_group state is INIT, RESYNC, or RECOVERY. To check the drive_group state, issue the following command:

```
$ sudo graidctl list drive_group
```

OR

```
$ sudo graidctl ls dg
```

Output Content

| Flag | Description |
|------------|---|
| DG ID | Drive group ID |
| MODE | Drive group RAID mode |
| VD NUM | Number of virtual drives in the drive group |
| CAPACITY | Total usable capacity of the drive group |
| FREE | Unused space of the drive group |
| USED | Used space of the drive group |
| CONTROLLER | Drive group controlled by the specific controller |
| STATE | Drive group state (see the following table) |

Drive Group State

| State | Description |
|--------------------|--|
| OFFLINE | Drive group is not working properly. This condition usually occurs when the number of damaged physical drives exceeds the limit. |
| OPTIMAL | Drive group is in optimal state. |
| OPTIMAL (!) | Drive group is in optimal state, but found inconsistency data. |
| OPTIMAL (cc) | Drive group is in optimal state and the consistency check task is ongoing. |
| OPTIMAL (cc!) | Drive group is in optimal state and the consistency check task is ongoing, but found inconsistent data. |
| DEGRADED | Drive group is available and ready, but the number of missing or failed physical drives has reached the limit. |
| PARTIALLY_DEGRADED | Drive group is available and ready for use, but some physical drives are missing or failed. |
| RECOVERY | Drive group is recovering |

| State | Description |
|--------|--|
| FAILED | Drive group is not working normally. |
| INIT | Drive group is initializing. |
| RESYNC | Drive group is resynchronizing. This condition usually occurs when the system encounters an abnormal crash. Do not replace the physical drive in this state until the resynchronization process completes. |
| RESCUE | Drive group is in rescue mode. |
| INIT | Drive group is initializing. |
| RESYNC | Drive group is resynchronizing. This condition usually occurs when the system encounters an abnormal crash. Do not replace the physical drive in this state until the re- synchronization process is complete. |

Deleting Drive Groups

To delete a drive group, issue the following command:

```
$ sudo graidctl delete drive_group <DG_ID> [flag]
```

OR

```
$ sudo graidctl del dg <DG_ID> [flag]
```

Note: You cannot delete a drive group that contains a virtual drive.

Output Example

In this example, drive group 1 was not deleted because it contains a virtual drive. Drive groups 0 and 2 were deleted successfully.

```
[graid@graid-demo ~]$ sudo graidctl del dg 1
✘Delete drive group failed: Failed to delete some DGs.
✘Delete drive group DG1 failed: rpc error: code = FailedPrecondition desc = DG1 still has 1VD(s)
[graid@graid-demo ~]$ sudo graidctl delete drive_group 0 2
✔Delete drive group DG0 successfully.
✔Delete drive group DG2 successfully.
```

Displaying Drive Group Information

To display detailed information about a drive group, issue the following command:

```
$ sudo graidctl describe drive_group <DG_ID> [flag]
```

OR

```
$ sudo graidctl desc dg <DG_ID> [flag]
```

| Flag | Description |
|------------|--|
| -h, --help | Help for the describe driver_group command |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl describe drive_group 0
✓Describe drive group successfully.
DG ID:                0
NQN:                  nqn.2020-05.com.graidtech:GRAID-SRD71ED4BAAAE12866
Model:                GRAID-SR
Serial:               D71ED4BAAAE12866
Firmware:             1.3.0-rc1
Mode:                 RAID5
Capacity:             10 TiB (11521889402880 B)
Free Space:           0 B
Used Space:           10 TiB (11521889402880 B)
Strip Size:           4096
State:                OPTIMAL
PD IDs:               [0 1 2 3]
Number of VDs:        1
Prefer Controller:    0
Running Controller:   0
Attributes:
    auto_failover = ENABLE
    rebuild_speed = high

[graid@graid-demo ~]$ sudo graidctl desc dg 1
✓Describe drive group successfully.
DG ID:                1
NQN:                  nqn.2020-05.com.graidtech:GRAID-SR05D4B25C1B5B3BCF
Model:                GRAID-SR
Serial:               05D4B25C1B5B3BCF
Firmware:             1.3.0-rc1
Mode:                 RAID10
Capacity:             40 GiB (42813358080 B)
Free Space:           40 GiB (42813358080 B)
Used Space:           0 B
Strip Size:           131072
State:                OPTIMAL
PD IDs:               [8 9 10 11]
Number of VDs:        1023
Prefer Controller:    0
Running Controller:   0
Attributes:
    rebuild_speed = high
    auto_failover = ENABLE
```

Output Content

| Flag | Description |
|--------------------|---|
| DG ID | Drive group ID |
| NQN | Drive group NQN |
| Model | Model number of the drive group |
| Serial | Serial number of the drive group |
| Firmware | Firmware version of the drive group |
| Mode | RAID mode of the drive group |
| Capacity | Capacity of the drive |
| Free Space | Remaining space on the drive |
| Used Space | Used space of the drive |
| Strip Size | Strip size (B) of the drive |
| PD IDs | All PDs of the drive |
| Number of VDs | Number of VDs of the drive Maximum: 1023 |
| Prefer Controller | Preferred controller of the drive |
| Running Controller | Running controller number of the drive |
| Attributes | Status of all attributes of the drive |

Selecting the Controller for a Drive Group

To set the controller to control a drive group, issue the following command:

```
$ sudo graidctl edit drive_group <DG_ID> controller <#>
```

Output Example

```
[graid@graid demo~]$ sudo graidctl list drive_group
✔List drive group successfully.
```

| DG ID | MODE | VD NUM | CAPACITY | FREE | USED | CONTROLLER | STATE |
|-------|-------|--------|----------|--------|---------|----------------------|---------|
| 0 | RAID1 | 1 | 3.5 TiB | 0 B | 3.5 TiB | running: 0 prefer: 0 | OPTIMAL |
| 1 | RAID5 | 35 | 10 TiB | 10 TiB | 0 B | running: 1 prefer: 1 | OPTIMAL |

```
[graid@graid demo~]$ sudo graidctl list drive_group
✔List controller successfully.
```

| ID | CONTROLLER MODEL | SERIAL NUMBER | NUMA | STATE | DG |
|----|------------------|---------------|------|--------|----|
| 0 | SR-1000 | 1xxxxxxxxxxx0 | 0 | ONLINE | 0 |
| 0 | SR-1000 | 1xxxxxxxxxxx1 | 1 | ONLINE | 1 |

```
[graid@graid demo~]$ sudo graidctl edit dg 1 controller 0
✔Edit drive group successfully.
[graid@graid demo~]$ sudo graidctl list drive_group
✔List drive group successfully.
```

| DG ID | MODE | VD NUM | CAPACITY | FREE | USED | CONTROLLER | STATE |
|-------|-------|--------|----------|--------|---------|----------------------|---------|
| 0 | RAID1 | 1 | 3.5 TiB | 0 B | 3.5 TiB | running: 0 prefer: 0 | OPTIMAL |
| 1 | RAID5 | 35 | 10 TiB | 10 TiB | 0 B | running: 0 prefer: 1 | OPTIMAL |

Assigning a Controller to a Drive Group

To assign a controller to control a drive group, issue the following command:

```
$ sudo graidctl create drive_group <RAID_Type> <PD_IDs> -c <#>
```

Output Example

```
[graid@graid demo~]$ sudo graidctl create drive_group raid5 0-3 -c 0
✔Create drive group successfully.
✔Create drive group DG0 successfully.
```

Managing Background Task Speed

To set the background task speed for a drive group, issue the following command:

```
$ sudo graidctl edit drive_group <DG_ID> rebuild_speed {low|normal|high}
```

Locating the Physical Drives in the Drive Group

To locate all the physical drives in a drive group, issue the following command:

```
$ sudo graidctl edit drive_group <DG_ID> locating start
```

To stop locating all the physical drives in a drive group, issue the following command:

```
$ sudo graidctl edit drive_group <DG_ID> locating stop
```

Degradation and Recovery

If multiple drive groups require simultaneous recovery, the drive groups recover individually.

If multiple physical drives in the same drive group require rebuilding, the physical drives are rebuilt simultaneously.

Rescue Mode

If a damaged drive group is initialized or a recovering drive group encounters an abnormal system crash, the data integrity of the drive group is affected. In this event, the drive group is forced offline to prevent data from being written to the drive group. To read the data for the drive group, force the drive group to go online using Rescue mode.

Note: A drive group in Rescue mode is read-only. Rescue mode cannot be disabled.

To enter rescue mode, issue the following command:

```
$ sudo graidctl edit drive_group <DG_ID> rescue_mode on
```

Managing Virtual Drives

Creating a Virtual Drive

To create a virtual drive, issue the following command:

```
$ sudo graidctl create virtual_drive <DG_ID> [<VD_SIZE>] [flags]
```

OR

```
$ sudo graidctl c vd <DG_ID> [<VD_SIZE>] [flags]
```

| Flag | Description |
|--------------|---|
| -h, --help | Help for the create virtual_drive command |
| -s, --serial | [string] Use user-specified serial ID |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl create virtual_drive 0
✓Create virtual drive VD0 in DG0 successfully.
[graid@graid-demo ~]$ sudo graidctl create virtual_drive 1 100G
✓Create virtual drive VD0 in DG1 successfully.
[graid@graid-demo ~]$ sudo graidctl create virtual_drive 2 1T
✓Create virtual drive VD0 in DG2 successfully.
```

Note: See [Setting Up the Auto-mount File Systems on Linux Using the SupremeRAID™ Driver](#) on page 136. It is critically important to follow these instructions to guarantee that the RAID group mounts automatically during system boot and to avoid any improper or unclear shutdown processes that could cause the RAID group to enter resync mode.

Listing Virtual Drives

To list virtual drives, issue the following command:

```
$ sudo graidctl list virtual_drive [flag]
```

OR

```
$ sudo graidctl ls vd [flag]
```

| Flag | Description |
|-------------|---|
| -h, --help | Help for the list virtual_drive command |
| -d, --dg-id | [string] List VDs of a certain DG ID |
| -v, --vd-id | [string] List certain VD IDs |

Output Example

```
root@graid:/home/graid# graidctl list virtual_drive
✔List virtual drive successfully.
```

| VD ID | DG ID | SIZE | DEVICE PATH | STATE | EXPORTED |
|-------|-------|---------|-------------|---------|----------|
| 0 | 0 | 959 MiB | /dev/gdg0n1 | OPTIMAL | No |

Output Content

| Flag | Description |
|-------------|--|
| DG ID | Drive group ID |
| VD ID | Virtual drive ID |
| SIZE | Usable size of the virtual drive |
| DEVICE PATH | Device path of the virtual drive |
| NQN | NQN of the virtual drive |
| STATE | Virtual drive state - identical to the drive group state (see the following table) |

| Flag | Description |
|----------|---|
| EXPORTED | Shows whether the virtual drive was exported using NVMe-oF or iSCSI |

Note: Do not perform I/O before the virtual drive is initialized and the device path (for example, /dev/gdgXnY) is created.

Virtual Drive State

| State | Description |
|--------------------|--|
| OFFLINE | Drive group is not working normally. This condition is usually caused when the number of damaged physical drives exceeds the limit. |
| OPTIMAL | Drive group is in the optimal state. |
| PARTIALLY_DEGRADED | Drive group is available and ready for use, but some physical drives are missing or failed. |
| RECOVERY | Drive group is recovering. |
| FAILED | Drive group is not working normally. |
| INIT | Drive group is initializing. |
| RESYNC | Drive group is resynchronizing. This condition usually occurs when the system encounters an abnormal crash. Do not replace the physical drive in this state until the resynchronization process completes. |
| RESCUE | Drive group is in rescue mode. |

Stripe Cache State

| State | Description |
|---------|--|
| OFFLINE | Stripe cache drive group is OFFLINE. |
| CLEAN | Stripe cache write-back has finished. |
| PURGE | Stripe cache is writing data into the virtual drive. |
| ACTIVE | Stripe cache is in optimal state. |

Deleting Virtual Drives

To delete virtual drives, issue the following command:

```
$ sudo graidctl delete virtual_drive <DG_ID> <VD_ID> [flags]
```

OR

```
$ sudo graidctl del vd <DG_ID> <VD_ID> [flags]
```

| Flag | Description |
|-------------|---|
| -h, --help | Help for the delete virtual_drive command |
| -f, --force | Delete VD forcibly |

Output Example

The following example shows that a virtual drive being used by the application cannot be deleted without adding the force flag.

```
[graid@graid-demo ~]$ sudo graidctl delete virtual_drive 0 0
✓Delete virtual drive VD0 from DG0 successfully.
[graid@graid-demo ~]$ sudo graidctl delete virtual_drive 2 0-1
✓Delete virtual drive VD1 from DG2 successfully.
✓Delete virtual drive VD0 from DG2 successfully.
```


Displaying Virtual Drive Information

To display detailed information about a virtual drive, issue the following command:

```
$ sudo graidctl describe virtual_drive <DG_ID> <VD_ID> [flags]
```

OR

```
$ sudo graidctl desc vd <DG_ID> <VD_ID> [flags]
```

| Flag | Description |
|------------|---|
| -h, --help | Help for the describe virtual_drive command |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl describe virtual_drive 0 4
✔Describe virtual drive successfully.
DG ID:          0
VD ID:          0
Serial:         EBFBC79373ED375F
DevicePath:     /dev/gdg0n1
Size:           4.3 GB
State:          OPTIMAL
Description:
Exported:
┌───┬───┬───┬───┬───┬───┐
│ PORT │ TRANSPORT TYPE │ ADDRESS │ INTERFACE │ ADDRESS FAMILY │ SERVICE ID │
├───┬───┬───┬───┬───┬───┤
│ 0 │ tcp │ 172.16.11.64 │ ens192 │ ipv4 │ 4420 │
└───┬───┬───┬───┬───┬───┘
```

Setting Up a Stripe Cache

Setting up a stripe cache improves HDD RAID 5 and RAID 6 sequential write performance. To set up a stripe cache:

Step 1 Create a stripe cache with a 4GB virtual drive.

```
$ sudo graidctl create virtual_drive 0 4GB
```

Note: For best practices, use a 4GB stripe whenever possible.

Step 2 Assign a 4GB virtual disk as the stripe cache.

```
$ sudo graidctl edit virtual_drive 0 0 stripecache 1 0
```

Step 3 Check the stripe cache.

```
$ sudo graidctl list virtual_drive
```

To flush the stripe cache, issue the following command:

```
$ sudo graidctl edit vd 0 0 stripecache none
```

Output Example

The assigned virtual drive is listed as = **Stripe Cache** = in the **DEVICE PATH** column.

```
[graid@graid-sake ~]$ sudo graidctl create virtual_drive 0
✓Create virtual drive successfully.
✓Create virtual drive DG0/VDO successfully.
[graid@graid-sake ~]$ sudo graidctl create virtual_drive 1 4GB
Create virtual drive successfully.
✓Create virtual drive DG1/VDO successfully.
[graid@graid-sake ~]$ sudo graidctl edit virtual_drive 0 0 stripecache 0 1
✓Edit virtual drive successfully.
[graid@graid-sake ~]$ sudo graidctl list virtual_drive
✓List virtual drive successfully.
```

| VD ID (2) | DG ID | SIZE | DEVICE PATH | STATE | EXPORTED |
|-----------|-------|--------|------------------|----------------------|----------|
| 0 | 0 | 9.3 GB | /dev/gdg0n1 | OPTIMAL cache:ACTIVE | No |
| 0 | 1 | 4.0 GB | Cache of DG0 VD0 | OPTIMAL | No |

Managing Controllers

Activating a Controller

To enable a controller, issue the following command:

```
$ sudo graidctl enable controller <Controller_ID> [flags]
```

OR

```
$ sudo graidctl enable cx <Controller_ID> [flags]
```

| Flag | Description |
|------------|--|
| -h, --help | Help for the enable controller command |

Output Example

```
[graid@graid demo~]$ sudo graidctl enable controller 0
✓Enable controller successfully.
✓Enable controller Controller 0 successfully.
[graid@graid demo~]$ sudo graidctl enable cx 1
✗Enable controller failed: Not found controller 1
```

Deactivating a Controller

To disable a controller, issue the following command:

```
$ sudo graidctl disable controller <Controller_ID> [flags]
```

OR

```
$ sudo graidctl disable cx <Controller_ID> [flags]
```

| Flag | Description |
|------------|---|
| -h, --help | Help for the disable controller command |

Output Example

```
[graid@graid demo~]$ sudo graidctl disable controller 0
✓Disable controller successfully.
✓Disable controller Controller 0 successfully.
[graid@graid demo~]$ sudo graidctl disable cx 1
✗Disable controller failed: Not found controller 1
```

Listing Controllers

To list controllers, issue the following command:

```
$ sudo graidctl list controller [flag]
```

OR

```
$ sudo graidctl ls cx [flag]
```

| Flag | Description |
|------------|--------------------------------------|
| -h, --help | Help for the list controller command |

Output Example

```
[graid@graid demo~]$ sudo graidctl list controller
✓List controller successfully.
```

| ID | CONTROLLER MODEL | SERIAL NUMBER | NUMA | STATE | DG |
|----|------------------|---------------|------|---------|-----|
| 0 | SR-1000 | 1xxxxxxxxxxx0 | 0 | ONLINE | 0,1 |
| 1 | SR-1000 | 1xxxxxxxxxxx1 | 1 | OFFLINE | 2,3 |

```
[graid@graid demo~]$ sudo graidctl ls cx
✓List controller successfully.
```

| ID | CONTROLLER MODEL | SERIAL NUMBER | NUMA | STATE | DG |
|----|------------------|---------------|------|---------|-----|
| 0 | SR-1000 | 1xxxxxxxxxxx0 | 0 | ONLINE | 0,1 |
| 1 | SR-1000 | 1xxxxxxxxxxx1 | 1 | OFFLINE | 2,3 |

Deleting a Controller

To delete a controller, issue the following command:

```
$ sudo graidctl delete controller [flag]
```

OR

```
$ sudo graidctl del cx [flag]
```

| Flag | Description |
|------------|--|
| -h, --help | Help for the delete controller command |

Note: You must disable the SupremeRAID™ controller before you can delete it. Disabling the controller prevents further access to it and its associated drives, allowing you to delete the controller safely without affecting the system's operation.

Output Example

```
[graid@graid demo~]$ graidctl delete controller 1
✖Delete controller failed: Controller 1 is still online, please disable it first
[graid@graid demo~]$sudo graidctl disable controller 1
✔Disable controller successfully.
✔Disable controller Controller 1 successfully.
[graid@graid demo~]$ sudo graidctl delete controller 1
✔Delete controller successfully.
✔Delete controller Controller 1 successfully.
```

Replacing a Controller License Key

To replace a controller's license key, issue the following command:

```
$ sudo graidctl replace controller <Controller_ID> <License_Key> [flags]
```

OR

```
$ sudo graidctl en cx <Controller_ID> <License_Key> [flags]
```

| Flag | Description |
|------------|---|
| -h, --help | Help for the replace controller license key command |

Observe the following guidelines when replacing a controller license key:

- To replace the license key for a controller in SupremeRAID, disable the controller first to ensure that the controller is not in use and can be updated safely. Disabling the controller prevents further access to it or its associated drives, allowing you to safely replace the license key without affecting the operation of the system.
- You cannot replace a license key with one that has a different architecture or supported features. Use the same license key or a compatible replacement to avoid replacement issues.
- If you are replacing a card in the system, deleting any inactive or invalid licenses associated with the old card is essential. Failing to do so may prevent other cards from becoming active, which is crucial for multi-controller systems.

Output Example

```
[graid@graid demo~]$ sudo graidctl replace controller 1 XXXXXXXX-XXXXXXX-XXXXXXX-XXXXXXX
✖Replace controller failed: Cannot replace ONLINE Controller 1
[graid@graid demo~]$sudo graidctl disable controller 1
✔Disable controller successfully.
✔Disable controller Controller 1 successfully.
[graid@graid demo~]$ sudo graidctl en cx 1 XXXXXXXX-XXXXXXX-XXXXXXX-XXXXXXX
✔Replace controller successfully.
✔Replace controller Controller 1 successfully.
```

Importing and Controlling MD Bootable NVMe RAIDs

After installing the SupremeRAID™ driver and the graidctl utility, SupremeRAID™ can import and control an MD bootable NVMe RAID. This feature makes it easy to swap drives if a bootable drive malfunctions.

Note: For instructions on setting up the MD bootable NVMe RAID, see [Configuring Boot-Drive Devices](#) on page 116.

Importing an MD Bootable NVMe RAID

Note: You can import only MD bootable NVMe RAID1.

To import an MD bootable NVMe RAID, issue the following command:

```
$ sudo graidctl import md_drive <DEVICE_PATH_0> <DEVICE_PATH_1> [flags]
```

OR

```
$ sudo graidctl imp md <DEVICE_PATH_0> <DEVICE_PATH_1> [flags]
```

| Flag | Description |
|------------|--------------------------------------|
| -h, --help | Help for the import md_drive command |

Output Example

```
[graid@graid ~]$ sudo graidctl import md_drive /dev/nvme0n1 /dev/nvme1n1
✓ Import md drive Import MD drives /dev/nvme0n1 /dev/nvme1n1 successfully.
[graid@graid ~]$ sudo graidctl ls pd
✓ List physical drive successfully.
```

| PD ID | DG ID | MQN/WWID | MODEL | CAPACITY | SLOT ID | STATE |
|-------|-------|--|--------------------------|----------|---------|--------|
| 32 | 4 | nqn.2014-08.org.nvmexpress:uuid:527978f1-8f0f-27b3-fb2f-8462d3c8f972 | VMware Virtual NVMe Disk | 27 GB | N/A | ONLINE |
| 33 | 4 | nqn.2014-08.org.nvmexpress:uuid:5218a65c-e259-6392-ff5c-35759b31b537 | VMware Virtual NVMe Disk | 27 GB | N/A | ONLINE |

```
[graid@graid ~]$ sudo graidctl ls dg
✓ List drive group successfully.
```

| DG ID | MODE | VD NUM | CAPACITY | FREE | USED | STATE |
|-------|-------|--------|----------|------|-------|---------|
| 4 | RAID1 | 3 | 27 GB | 0 B | 27 GB | OPTIMAL |

```
[graid@graid ~]$ sudo graidctl ls vd
✓ List virtual drive successfully.
```

| VD ID | DG ID | SIZE | DEVICE PATH | STATE |
|-------|-------|--------|-------------|---------|
| 0 | 4 | 11 GB | /dev/md127 | OPTIMAL |
| 1 | 4 | 5.4 GB | /dev/md125 | OPTIMAL |
| 2 | 4 | 5.4 GB | /dev/md126 | OPTIMAL |

Replacing an MD Bootable NVMe RAID1

Note: You can replace only MD bootable NVMe RAID1.

To replace an MD bootable NVMe RAID 1, replace the old NVMe SSD with the new one. The old physical drive state should indicate **MISSING**.

```
$ sudo graidctl replace md_drive <OLD_MD> PD_ID> <NEW_DEVICE_PATH> [flags]
```

OR

```
$ sudo graidctl en md <OLD_MD> PD_ID> <NEW_DEVICE_PATH> [flags]
```

| Flag | Description |
|-------------|---------------------------------------|
| -h, --help | Help for the replace md_drive command |
| -f, --force | Replace ONLINE MD forcibly |

Output Example

The following example shows an MD missing.

```
[graid@graid ~]$ sudo graidctl ls pd
✓ List physical drive successfully.
```

| PD ID | DG ID | NQN/WWID | MODEL | CAPACITY | SLOT ID | STATE |
|-------|-------|---|--------------------------|----------|---------|---------|
| 32 | 4 | nqn.2014-08.org.nvmeexpress:uuid:527970f1-8f0f-27b3-fb2f-8462d3c8f972 | VMware Virtual NVMe Disk | 27 GB | N/A | ONLINE |
| 33 | 4 | nqn.2014-08.org.nvmeexpress:uuid:5218a65c-e259-6392-ff5c-35759b31b537 | VMware Virtual NVMe Disk | 27 GB | N/A | MISSING |

```
[graid@graid ~]$ sudo graidctl ls dg
✓ List drive group successfully.
```

| DG ID | MODE | VD NUM | CAPACITY | FREE | USED | STATE |
|-------|-------|--------|----------|------|-------|----------|
| 4 | RAID1 | 3 | 27 GB | 0 B | 27 GB | DEGRADED |

```
[graid@graid ~]$ sudo graidctl ls vd
✓ List virtual drive successfully.
```

| VD ID | DG ID | SIZE | DEVICE PATH | STATE |
|-------|-------|--------|-------------|----------|
| 0 | 4 | 11 GB | /dev/md127 | DEGRADED |
| 1 | 4 | 5.4 GB | /dev/md125 | DEGRADED |
| 2 | 4 | 5.4 GB | /dev/md126 | DEGRADED |

The following example shows a replaced drive. The bootable RAID group rebuilds immediately after replacing the drive.

```
[graid@graid ~]$ sudo graidctl replace md_drive 33 /dev/nvme2n1
✓ Replace md drive replaced PD33 with /dev/nvme2n1 successfully.
[graid@graid ~]$ sudo graidctl ls pd
✓ List physical drive successfully.
```

| PD ID | DG ID | NQN/WWID | MODEL | CAPACITY | SLOT ID | STATE |
|-------|-------|---|--------------------------|----------|---------|---------|
| 32 | 4 | nqn.2014-08.org.nvmeexpress:uuid:527970f1-8f0f-27b3-fb2f-8462d3c8f972 | VMware Virtual NVMe Disk | 27 GB | N/A | ONLINE |
| 33 | 4 | nqn.2014-08.org.nvmeexpress:uuid:52524729-5a31-13e7-a316-f6e765e16ec8 | VMware Virtual NVMe Disk | 27 GB | N/A | REBUILD |

```
[graid@graid ~]$ sudo graidctl ls dg
✓ List drive group successfully.
```

| DG ID | MODE | VD NUM | CAPACITY | FREE | USED | STATE |
|-------|-------|--------|----------|------|-------|---------|
| 4 | RAID1 | 3 | 27 GB | 0 B | 27 GB | REBUILD |

```
[graid@graid ~]$ sudo graidctl ls vd
✓ List virtual drive successfully.
```

| VD ID | DG ID | SIZE | DEVICE PATH | STATE |
|-------|-------|--------|-------------|-------------------|
| 0 | 4 | 11 GB | /dev/md127 | REBUILD (pending) |
| 1 | 4 | 5.4 GB | /dev/md125 | REBUILD (82.52%) |
| 2 | 4 | 5.4 GB | /dev/md126 | REBUILD (pending) |

Dismissing an Imported MD Bootable NVMe RAID1

Note: You can dismiss only MD bootable NVMe RAID1.

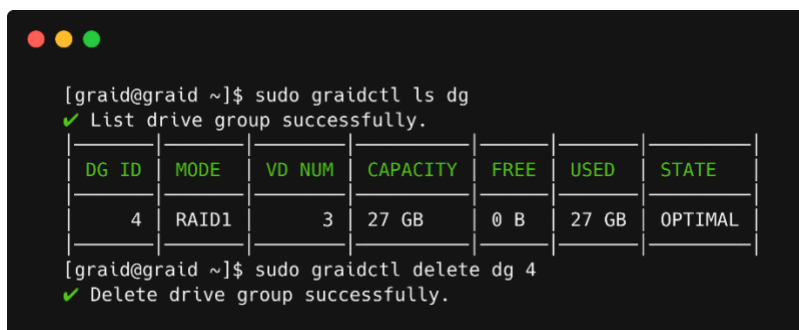
To dismiss an imported MD bootable NVMe RAID 1, issue the following command:

```
$ sudo graidctl delete drive_group <DG_ID> [flags]
```

OR

```
$ sudo graidctl del dg <DG_ID> [flags]
```

Output Example



```
[graid@graid ~]$ sudo graidctl ls dg
✓ List drive group successfully.
```

| DG ID | MODE | VD NUM | CAPACITY | FREE | USED | STATE |
|-------|-------|--------|----------|------|-------|---------|
| 4 | RAID1 | 3 | 27 GB | 0 B | 27 GB | OPTIMAL |

```
[graid@graid ~]$ sudo graidctl delete dg 4
✓ Delete drive group successfully.
```

Adjusting or Updating Configuration Settings for the SupremeRAID™ Add-on

The add-on for SupremeRAID™ provides enhanced configuration options and allows you to fine-tune system settings to meet your specific needs. Follow these steps to ensure that the add-on is configured optimally for maximum system performance.

Editing Configuration Settings

To edit the configuration, issue the following command:

```
$ sudo graidctl edit config <config_name> <value> [flags]
```

OR

```
$ sudo graidctl e conf <config_name> <value> [flags]
```

| Flag | Description |
|------------|----------------------------------|
| -h, --help | Help for the edit config command |

Configuration Options

| Field | Description |
|-------|--|
| LED | Importing LED configuration files |
| SED | Add single SED key for specific device |

Output Example

```
[graid@graid demo~]$ sudo graidctl edit config sed nqn.2019-08.org.qemu:NVME0002
Enter Key: ✓Edit config successfully.
[graid@graid demo~]$ sudo graidctl e conf led R282-Z9G.yaml
✓Edit config successfully.
```

Describing Configuration Settings

To describe the configuration, issue the following command:

```
$ sudo graidctl describe config <config_name> [flags]
```

OR

```
$ sudo graidctl desc conf <config_name> [flags]
```

| Flag | Description |
|------------|--------------------------------------|
| -h, --help | Help for the describe config command |

Configuration Options

| Field | Description |
|-------|---|
| LED | Obtain the imported LED configuration files |

| Field | Description |
|-------|--------------------------------|
| SED | Obtain the SED key information |

Output Example

```
[graid@graid demo~]$ sudo graidctl describe config sed
✔Describe config successfully.
Totally 1 SED keys.
Device GUIDs:
  nqn.2019-08.org.qemu:NVME0002
[graid@graid demo~]$ sudo graidctl desc conf led
✘Describe config failed: LED config does not exist
```

Deleting Configuration Settings

To delete the configuration, issue the following command:

```
$ sudo graidctl delete config <config_name> [flags]
```

OR

```
$ sudo graidctl del conf <config_name> [flags]
```

| Flag | Description |
|------------|------------------------------------|
| -h, --help | Help for the delete config command |

Configuration Options

| Field | Description |
|-------|---|
| LED | Obtain the imported LED configuration files |
| SED | Obtain the SED key information |

Output Example

```
[graid@graid demo~]$ sudo graidctl delete config sed all
Do you really want to delete all SED key?
Repeat IMEANTODELETEALL to continue: IMEANTODELETEALL
✔Delete config successfully.
[graid@graid demo~]$ sudo graidctl del conf led
✔Delete config successfully.
```

Restoring SupremeRAID™ Configuration Settings

To scan all NVMe and SCSI drives and restore the latest SupremeRAID™ configuration, issue the following command:

```
$ sudo graidctl restore config [flags]
```

OR

```
$ sudo graidctl re conf [flags]
```

| Flag | Description |
|------------|--|
| -h, --help | Help for the restore config command |
| -a, --auto | Selects the last configuration automatically |

Output Example

```
[graid@graid demo~]$ sudo graidctl restore config
✖Restore config failed: Please stop the graid service before restoring the config, and restart the graid service after restored the config.
[graid@graid demo~]$ sudo graidctl re conf
Skip /dev/sda: no config found
Found the following configs:
0: Device /dev/nvme0n1, UUID 00200000-0000-0000-4d02-000000000000, Epoch 1412, Time 2022-12-08 20:14:09 +0800 CST
1: Device /dev/nvme1n1, UUID 00200000-0000-0000-4d02-000000000000, Epoch 1412, Time 2022-12-08 20:14:09 +0800 CST
2: Device /dev/nvme2n1, UUID 00200000-0000-0000-4d02-000000000000, Epoch 1412, Time 2022-12-08 20:14:09 +0800 CST
3: Device /dev/nvme3n1, UUID 00200000-0000-0000-4d02-000000000000, Epoch 1412, Time 2022-12-08 20:14:09 +0800 CST
Please select one config to restore (0-3): 0
Restore to /etc/graid.conf (y/N)? y
✔Restore config graid.conf successfully.
```

Managing Events

Listing Events

To check detailed information from record, issue the following command:

```
$ sudo graidctl list event [flags]
```

OR

```
$ sudo graidctl ls event [flags]
```

| Flag | Description |
|-------------------|---|
| -h, --help | Help for the list event command |
| -c, --component | [string] Filter events by component |
| -n, --max_entries | [int32] Limit the number of events returned |
| -o, --output | [string] Output to a file |
| -s, --severity | [string] Filter events by severity |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl list event -n 10 -s INFO -c DG
✓List event successfully.
[2022-06-22 22:06:29 +0800 CST][INFO][DG][0] State transitted from UNKNOWN to OFFLINE.
[2022-06-22 22:20:07 +0800 CST][INFO][DG][0] Drive group deleted.
[2022-06-22 22:21:13 +0800 CST][INFO][DG][0] State transitted from UNKNOWN to OPTIMAL.
[2022-06-22 22:21:13 +0800 CST][INFO][DG][0] Drive group created.
[2022-06-22 22:28:02 +0800 CST][INFO][DG][0] Drive group deleted.
[2022-06-22 22:28:20 +0800 CST][INFO][DG][0] State transitted from UNKNOWN to OPTIMAL.
[2022-06-22 22:28:20 +0800 CST][INFO][DG][0] Drive group created.
[2022-06-22 22:30:15 +0800 CST][INFO][DG][0] CC has started.
[2022-06-22 23:26:57 +0800 CST][INFO][DG][0] CC has completed.
[2022-06-22 23:26:57 +0800 CST][INFO][DG][0] CC has started.
```

Deleting Events

To delete events, issue the following command:

```
$ sudo graidctl delete event [flags]
```

OR

```
$ sudo graidctl del event [flags]
```

| Flag | Description |
|---------------|---|
| -h, --help | Help for the delete event command |
| -d, --date | [string] Delete event entries before the date |
| -e, --entries | int32] Keep the latest number of entries Default: -1 |

Managing Remote NVMe-oF Targets

Before you can create physical drives from NVMe-oF devices, you must connect to the NVMe-oF target.

Connecting to a Remote NVMe-oF Target

To connect to a remote NVMe-oF target, issue the following command:

```
$ sudo graidctl connect remote_target <transport type> <addr> <address family> <port service id>
```

OR

```
$ sudo graidctl con rt <transport type> <addr> <address family> <port service id>
```

| Flag | Description |
|------------|--|
| -h, --help | Help for the connect remote_target command |

Required Parameters

| Option | Description |
|----------------|---|
| transport type | Network fabric used for a NVMe-over-Fabrics network. Current string values include: <ul style="list-style-type: none"> RDMA = network fabric is an RDMA network (RoCE, iWARP, InfiniBand, basic RDMA, etc.) TCP = network fabric is a TCP/IP network. |
| ip address | Network address of the controller |
| address family | Network address protocol. Current string values include ipv4/ipv6. |
| port service | Transport service ID |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl connect remote_target rdma 192.168.2.10 ipv4 4420
✓Connect remote target successfully.
✓Connect remote target Target 0 successfully.
[graid@graid-demo ~]$ sudo graidctl connect remote_target tcp 192.168.2.11 ipv4 4420
✓Connect remote target successfully.
✓Connect remote target Target 1 successfully.
```

Listing Connected Remote NVMe-oF Targets

To list all of the connected NVMe-oF targets, issue the following command:

```
$ sudo graidctl list remote_target
```

OR

```
$ sudo graidctl ls rt
```

| Flag | Description |
|------------|---|
| -h, --help | Help for the list remote_target command |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl list nvmeof_target
✓List nvmeof target successfully.
```

| PORT ID | TYPE | INTERFACE | ADDRESS | ADDRESS FAMILY | SERVICE ID | SUBSYSTEMS |
|---------|------|-----------|--------------|----------------|------------|---------------------------|
| 0 | tcp | ens160 | 172.16.11.81 | ipv4 | 4420 | DG0/VD0, DG0/VD1 |
| 1 | tcp | ens160 | 172.16.11.81 | ipv4 | 4421 | DG0/VD0, DG0/VD1, DG0/VD3 |

```
[graid@graid-demo ~]$ sudo graidctl ls nt
✓List nvmeof target successfully.
```

| PORT ID | TYPE | INTERFACE | ADDRESS | ADDRESS FAMILY | SERVICE ID | SUBSYSTEMS |
|---------|------|-----------|--------------|----------------|------------|---------------------------|
| 0 | tcp | ens160 | 172.16.11.81 | ipv4 | 4420 | DG0/VD0, DG0/VD1 |
| 1 | tcp | ens161 | 172.16.11.82 | ipv4 | 4420 | DG0/VD0, DG0/VD1, DG0/VD3 |

Disconnecting from Remote NVMe-oF Targets

To disconnect from an NVMe-oF target, issue the following command:

```
$ sudo graidctl disconnect remote_target <target id>
```

OR

```
$ sudo graidctl dis rt <target id>
```

| Flag | Description |
|------------|---|
| -h, --help | Help for the disconnect remote_target command |

Note: You cannot delete the target when there are physical drives created from the target.

Output Example

```
[graid@graid-demo ~]$ sudo graidctl disconnect remote_target 0
✔Disconnect remote target successfully.
✔Disconnect remote target Port 0 successfully.
[graid@graid-demo ~]$ sudo graidctl dis rt 1
✔Disconnect remote target successfully.
✔Disconnect remote target Port 1 successfully.
```

Exporting NVMe-oF Target Management

You can export the virtual drive to other initiators.

Creating the NVMe-oF Target Port Service

To create the NVMe-oF target port service, issue the following command:

```
$ sudo graidctl create nvmeof_target <tcp|rdma> <interface> <address family>
<svrcid> [flags]
```

OR

```
$ sudo graidctl c nt <tcp|rdma> <interface> <address family> <svrcid> [flags]
```

| Flag | Description |
|------------|--|
| -h, --help | Help for the create nvmeof_targets command |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl create nvmeof_target tcp ens160 ipv4 4420
✔Create nvmeof target successfully.
✔Create nvmeof target Port 0 successfully.
[graid@graid-demo ~]$ sudo graidctl create nvmeof_target tcp ens161 ipv4 4420
✔Create nvmeof target successfully.
✔Create nvmeof target Port 1 successfully.
```

Exporting NVMe-oF Targets

To export NVMe-oF targets using the service port you created, issue the following command:

```
$ sudo graidctl export virtual_drive <DG_ID> <VD_ID> [flags]
```

OR

```
$ sudo graidctl exp vd <DG_ID> <VD_ID> [flags]
```

| Flag | Description |
|----------------|---|
| -h, --help | Help for the export NVMe-oF targets command |
| -a, --all | Export all NVMe-oF target into all ports |
| -p, --port-ids | Port IDs [Int32] |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl export virtual_drive 0 0-1 --all
✓Export virtual drive successfully.
✓Export virtual drive VD0 into Port 0 successfully.
✓Export virtual drive successfully.
✓Export virtual drive VD1 into Port 0 successfully.
[graid@graid-demo ~]$ sudo graidctl export virtual_drive 0 3 --port-ids=1
✓Export virtual drive successfully.
✓Export virtual drive VD3 into Port 1 successfully.
[graid@graid-demo ~]$ sudo graidctl export vd 0 2 --port-ids=1
✓Export virtual drive successfully.
✓Export virtual drive VD2 into Port 1 successfully.
```

Listing Created NVMe-oF Targets

To list all created NVMe-oF target devices, issue the following command:

```
$ sudo graidctl list nvmeof_target
```

OR

```
$ sudo graidctl ls nt
```

| Flag | Description |
|------------|---|
| -h, --help | Help for the list nvmeof_target command |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl list nvmeof_target
✔List nvmeof target successfully.
```

| PORT ID | TYPE | INTERFACE | ADDRESS | ADDRESS FAMILY | SERVICE ID | SUBSYSTEMS |
|---------|------|-----------|--------------|----------------|------------|---------------------------|
| 0 | tcp | ens160 | 172.16.11.81 | ipv4 | 4420 | DG0/VD0, DG0/VD1 |
| 1 | tcp | ens160 | 172.16.11.81 | ipv4 | 4421 | DG0/VD0, DG0/VD1, DG0/VD3 |

```
[graid@graid-demo ~]$ sudo graidctl ls nt
✔List nvmeof target successfully.
```

| PORT ID | TYPE | INTERFACE | ADDRESS | ADDRESS FAMILY | SERVICE ID | SUBSYSTEMS |
|---------|------|-----------|--------------|----------------|------------|---------------------------|
| 0 | tcp | ens160 | 172.16.11.81 | ipv4 | 4420 | DG0/VD0, DG0/VD1 |
| 1 | tcp | ens161 | 172.16.11.82 | ipv4 | 4420 | DG0/VD0, DG0/VD1, DG0/VD3 |

Deleting the NVMe-oF Target Port Service Unexporting NVMe-oF Targets

To delete the NVMe-oF target port service, issue the following command:

```
$ sudo graidctl delete nvmeof_target <PORT_ID> [flags]
```

OR

```
$ sudo graidctl del nt <PORT_ID> [flags]
```

| Flag | Description |
|-------------|---|
| -h, --help | Help for the delete nvmeof_target command |
| -f, --force | Force delete ports |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl delete nvmeof_target 0
✓Delete nvmeof target successfully.
✓Delete nvmeof target Port 0 successfully.
[graid@graid-demo ~]$ sudo graidctl del nt 1
✓Delete nvmeof target successfully.
✓Delete nvmeof target Port 1 successfully.
```

Unexporting NVMe-oF Targets

To unexport an NVMe-oF target, issue the following command:

```
$ sudo graidctl unexport virtual_drive <DG_ID> <VD_ID> [flags]
```

OR

```
$ sudo graidctl unexp vd <DG_ID> <VD_ID> [flags]
```

| Flag | Description |
|------------|---|
| -h, --help | Help for the unexport nvmeof_target command |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl unexport virtual_drive 0 3 -p 1
✓Unexport virtual drive successfully.
✓Unexport virtual drive VD3 from Port 1 successfully.
[graid@graid-demo ~]$ sudo graidctl unexp vd 0 0-1 --all
✓Unexport virtual drive successfully.
✓Unexport virtual drive VD0 from Port 0 successfully.
✓Unexport virtual drive VD0 from Port 1 successfully.
✓Unexport virtual drive successfully.
✓Unexport virtual drive VD1 from Port 0 successfully.
✓Unexport virtual drive VD1 from Port 1 successfully.
```

Using Consistency Checks to Ensure Data Integrity

The consistency check operation verifies that the data is correct in DGs that use RAID levels 1, 5, 6, and 10. In a system with parity, for example, checking consistency calculates the data on one drive and compares the results to the contents of the parity drive.

Note: You cannot perform a consistency check on RAID 0 because it does not provide data redundancy. Additionally, a consistency check can only run when the DG is in `OPTIMAL` or `PARTIALLY_DEGRADED` state.

The consistency check function records all events to the event database, and `graidctl` provides commands to retrieve the events. The maximum number of event entries is 1,000. The system deletes event entries periodically. You can also delete entries manually.

Starting Consistency Checks Manually

To start a consistency check manually, issue the following command:

```
$ sudo graidctl start consistency_check manual_task [flags]
```

OR

```
$ sudo graidctl start cc [flags]
```

| Flag | Description |
|--------------|--|
| -h, --help | Help for the start consistency_check manual command |
| -p, --policy | [string] Specify CC policy [stop_on_error/auto_fix] |

DG State

Enabling a consistency check task adds the following annotations beside the output string of the DG state.

| DG State | Description |
|---------------|---|
| OPTIMAL | Normal state without enabling consistency check |
| OPTIMAL (!) | Inconsistency found |
| OPTIMAL (cc) | Consistency check ongoing |
| OPTIMAL (cc!) | Consistency check ongoing and inconsistency found |

Output Example



```
[graid@graid-demo ~]$ sudo graidctl start consistency_check manual_task 0 1 -p stop_on_error
✓Start consistency check successfully.
[graid@graid-demo ~]$ sudo graidctl start cc manual_task 2 -p auto_fix
✓Start consistency check successfully.
```

Stopping Consistency Check

To stop a consistency check task, issue the following command:

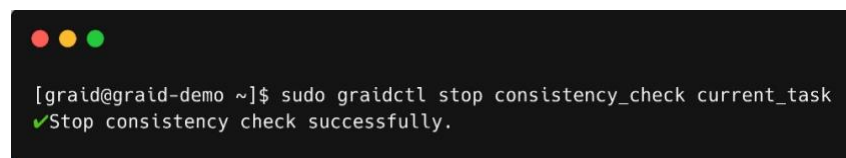
```
$ sudo graidctl stop consistency_check current_task [flags]
```

OR

```
$ sudo graidctl stop cc current_task [flags]
```

| Flag | Description |
|------------|---|
| -h, --help | Help for the stop consistency_check command |

Output Example



```
[graid@graid-demo ~]$ sudo graidctl stop consistency_check current_task
✓Stop consistency check successfully.
```

Scheduling Consistency Checks

To schedule a consistency check task, issue the following command:

```
$ sudo graidctl set consistency_check schedule_mode
<off|continuously|hourly|daily|weekly|monthly><yyyy/mm/dd> <hh> [flags]
```

OR

```
$ sudo graidctl set cc schedule_mode
<off|continuously|hourly|daily|weekly|monthly> <yyyy/mm/dd> <hh> [flags]
```

| Flag | Description |
|------------|--|
| -h, --help | Help for the set consistency_check command |

DG State

Enabling a consistency check task adds the following annotations beside the output string of the DG state.

| DG State | Description |
|---------------|---|
| OPTIMAL | Normal state without enabling consistency check |
| OPTIMAL (!) | Inconsistency found |
| OPTIMAL (cc) | Consistency check ongoing |
| OPTIMAL (cc!) | Consistency check ongoing and inconsistency found |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl set consistency_check schedule_mode daily 2022/06/25 10
✔Set consistency check successfully.
```

Viewing Consistency Check Information

To view detailed consistency check information, issue the following command:

```
$ sudo graidctl describe consistency_check [flags]
```

OR

```
$ sudo graidctl desc consistency_check [flags]
```

| Flag | Description |
|------------|---|
| -h, --help | Help for the describe consistency_check command |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl describe consistency_check
✓Describe consistency check successfully.
Schedule Mode:      daily
Schedule Base:     2022-06-25 10:00:00 +0800 CST
Excluded DGs:      []
Policy:            stop_on_error
Next Schedule:    2022-06-26 10:00:00 +0800 CST
Current Task:     2 DG(s)
                  -DG0: Checking (progress: 28.15%)
                     Start Time: 2022-06-26 09:37:37 +0800 CST
                     End Time:
                  -DG1: Pending
                     Start Time:
                     End Time:
```

Setting the Consistency Check Policy

To set a consistency check policy, issue the following command.

Note: By default, the consistency check runs on all drive_groups. To exclude drive groups, run the `xcluded_dgs` command.

```
$ sudo graidctl set consistency_check policy <auto_fix|stop_on_error> [flags]
```

| Flag | Description |
|------------|--|
| -h, --help | Help for the set consistency_check command |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl set consistency_check policy auto_fix
✓Set consistency check successfully.
```

Excluding Drive Groups from the Consistency Check Policy

To exclude some drive groups from a consistency check policy, issue the following command:

```
$ sudo graidctl set consistency_check excluded_dgs <DG_IDs>
```

OR

```
$ sudo graidctl set cc excluded_dgs <DG_IDs>
```

| Flag | Description |
|------------|--|
| -h, --help | Help for the set consistency_check command |

Output Example

```
[graid@graid-demo ~]$ sudo graidctl set consistency_check excluded_dgs 1
✔Set consistency check successfully.
```

ADDITIONAL FUNCTIONS

This chapter describes the following additional tasks you can perform with SupremeRAID™.

- Configuring Boot-Drive Devices
- Manually Migrating the RAID Configuration Between Hosts
- Restarting the SupremeRAID™ Service After Upgrading the System Kernel
- Obtaining SMART Information from Devices
- Monitoring System Input/Output Statistics for Devices Using iostat
- Setting Up the Auto-mount File Systems on Linux Using the SupremeRAID™ Driver
- ESXi Virtual Machine Support Using GPU Passthrough
- Using Self-Encrypting Drives

Configuring Boot-Drive Devices

You can configure two NVMe SSDs as RAID1 boot devices and control them using SupremeRAID™. The procedure you use depends on the operating system.

- For CentOS, see [Procedure for CentOS on page 117](#).
- For SLES 15 SP2 and SP3, see [Procedure for SLES 15 SP2, and SP3 on page 125](#).

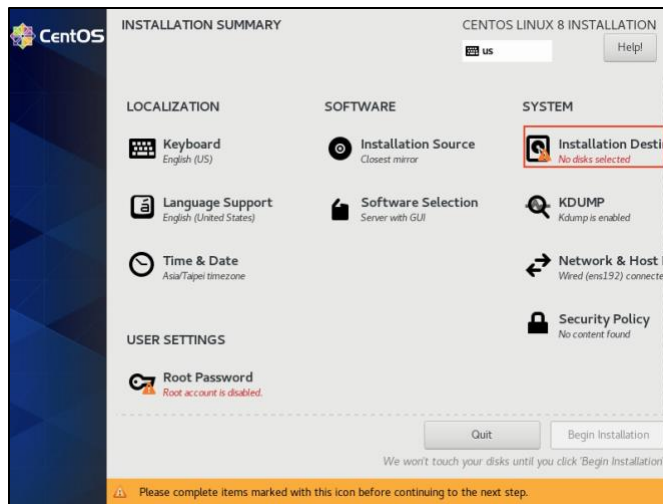
Note: You cannot configure boot-drive devices across multiple operating systems.

Procedure for CentOS

Assigning RAID1 Boot Devices Manually

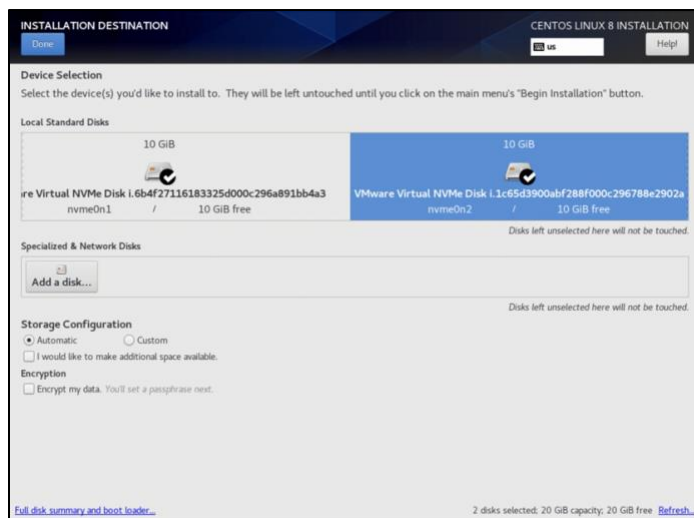
You assign RAID1 boot devices when you install CentOS. If the CentOS GUI does not prompt you to assign the boot devices, you can assign them manually.

Step 1 From the INSTALLATION SUMMARY page, select **SYSTEM > Installation Destination**.



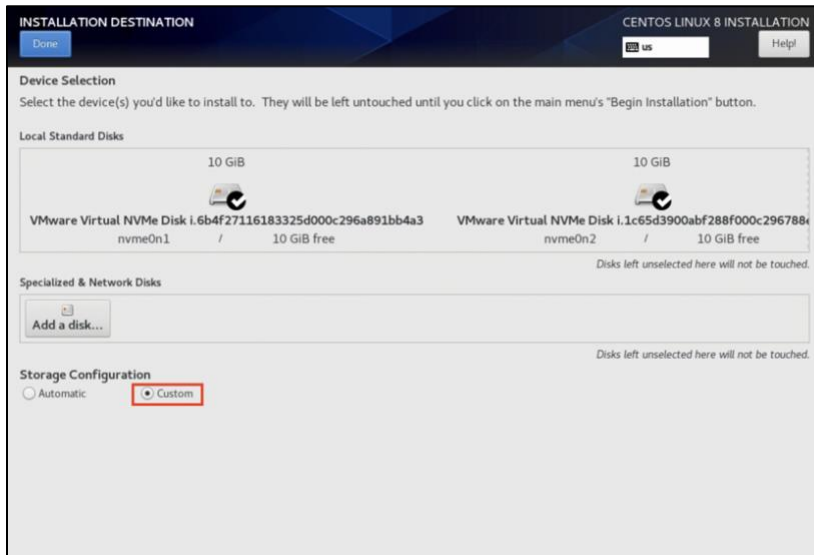
Step 2 From the INSTALLATION DESTINATION page, select the two NVMe SSDs that you want to set as RAID1 boot devices.

Tip: To select multiple devices, use the Ctrl key.



Step 3 For Storage Configuration, select Custom.

Step 4 Click Done.




Creating Storage Partitions Manually

You manually create the storage partitions on CentOS systems.


Step 1 From the MANUAL PARTITIONING page, select New CentOS Linux 8 Installation.

Step 2 Click here to create them automatically to create the mount points.



Step 3 Set Device Type to RAID and set RAID LEVEL to RAID 1.


Step 4 Click Update Settings.



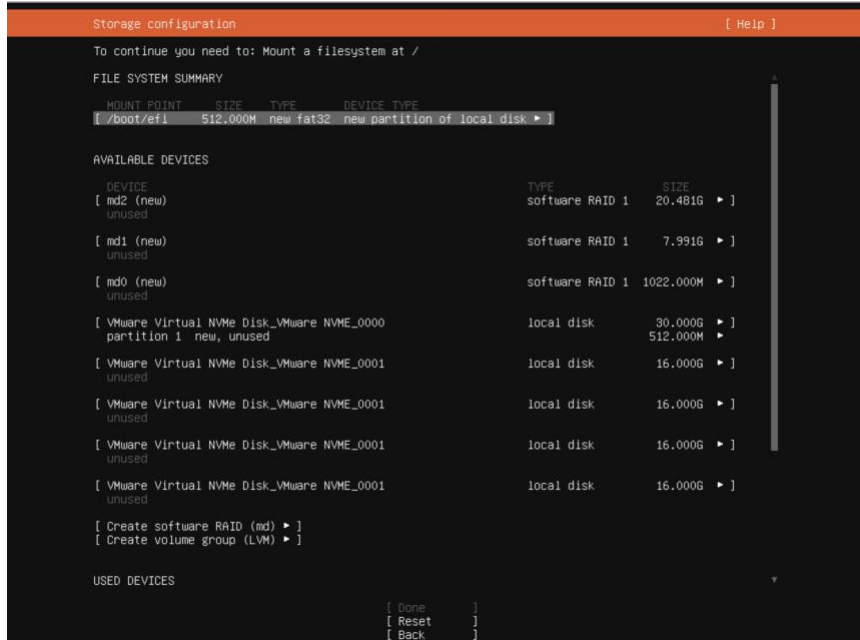
Creating and Configuring Storage Partitions

Storage partitions must be created and configured during the Ubuntu Server 20.04 installation. The partitions are required for mounting /boot, swap, and root/. Each partition functions as a soft RAID.

Step 1 From the Guided storage configuration page, select **Custom storage layout**.



Step 2 From the Storage configuration page, select the first disk as the boot disk.




Step 3 From the second Disk menu, select Add GPT partition > Create a partition.

Step 4 Set the size of the new partition. Use the same size as the boot disk so that the first and second partitions align.

Step 5 For Format, select [Leave unformatted].

Note: You must use [Leave unformatted]. DO NOT mount the partition. Setting RAID1 and mounting partitions on multiple drives (MD) occurs later in this procedure.

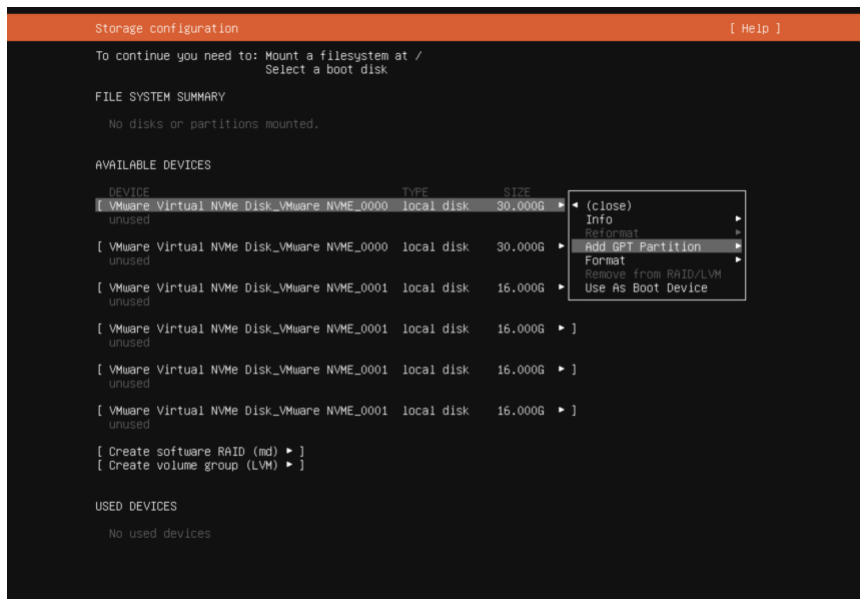
Step 6 Select [Create] to create the storage partition.



Configuring the Boot Partitions

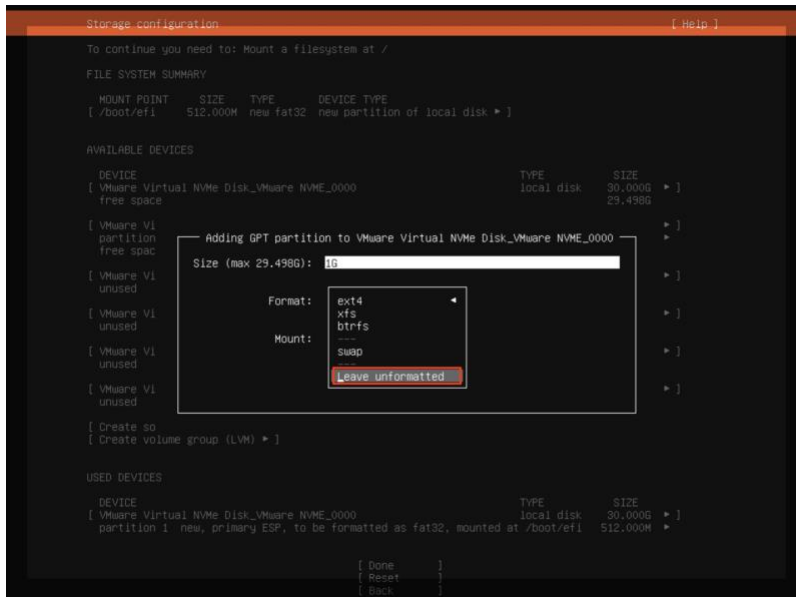
The following procedure describes how to configure the /boot, swap, and root/ partitions on both disks.

Step 1 From the Storage configuration page Disk menu, select **Add GPT Partition**.



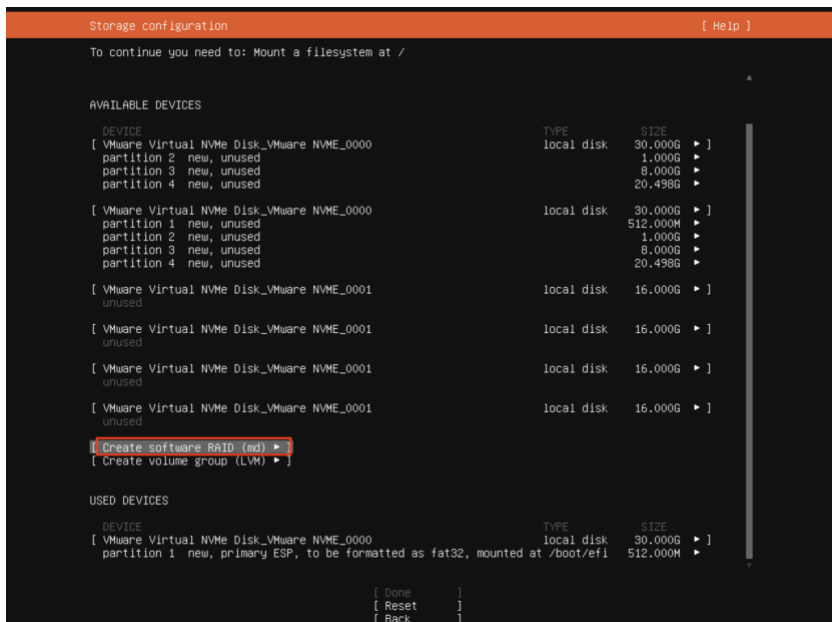
Step 2 Set the size of the partitions. Use 1G for /boot, the memory size for swap, and the remaining size for root/.

Step 3 For Format, select [Leave unformatted].



Creating a Software RAID for Multiple Devices (MD)

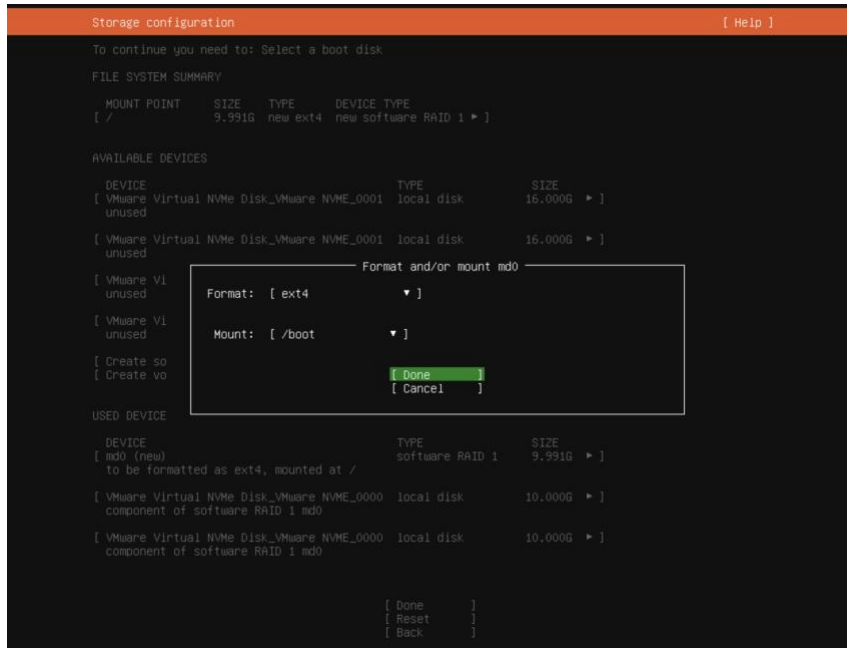
To create the software RAID on multiple devices, from the Storage configuration page, select **Create software RAID (md)**.



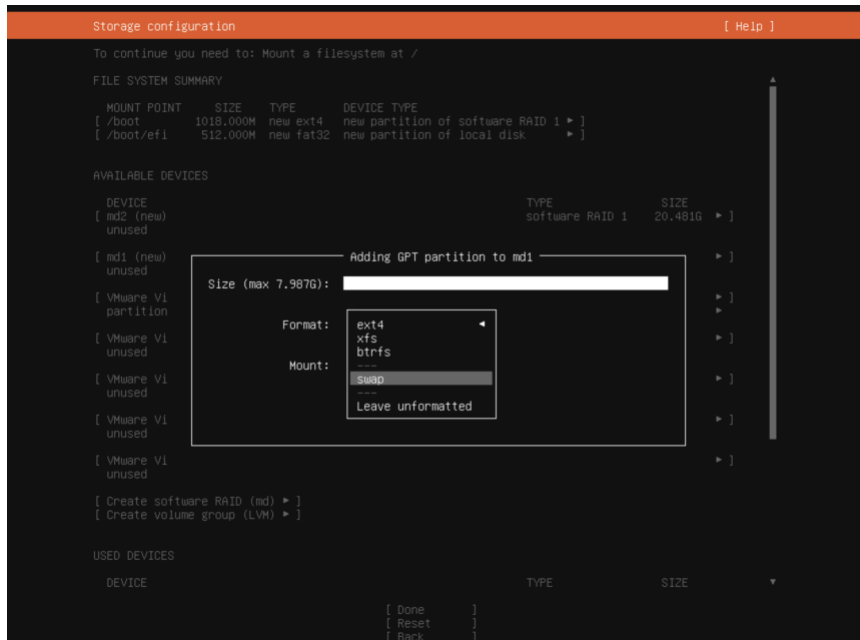
Setting MD as the Mounting Point

To set MD as the mounting point:


Step 1 From the Storage configuration page Disk menu, set md0 as the /boot mounting point.



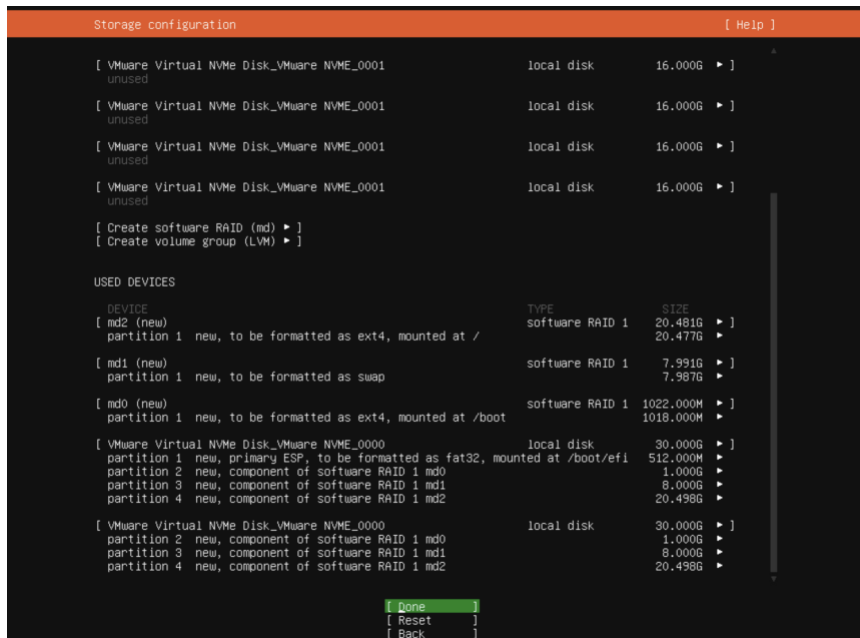
Step 2 From the Disk menu, select **Add GPT Partition** and set md1 as the swap mounting point.



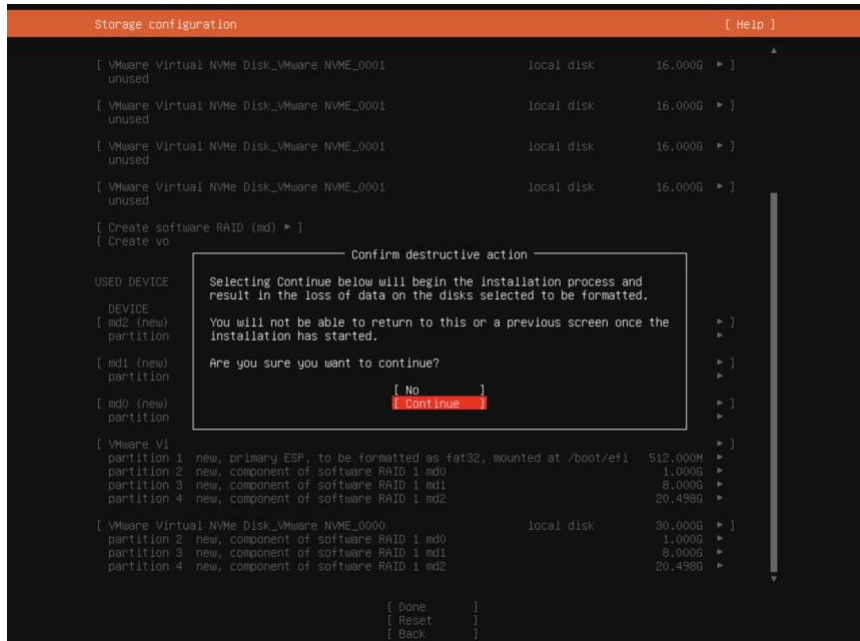
Step 3 From the **Disk** menu, select **Add GPT Partition** and set md2 as the root / mounting point.



Step 4 After setting the mount points, click **Done**.



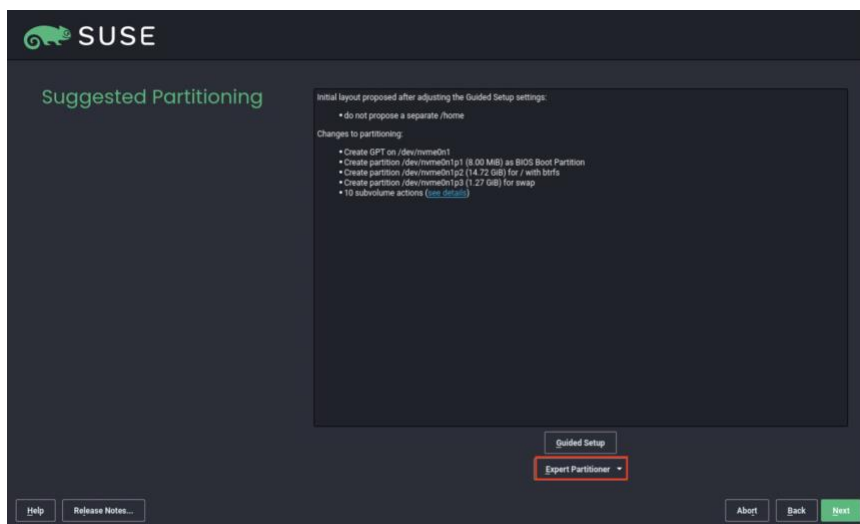
Step 5 From the Confirm destructive action popup, select **Continue**. The partition settings are now in effect.




Procedure for SLES 15 SP2, and SP3

When installing SLES 15 SP2 or SP3, you must manually create RAID1 and configure the partitions. To manually create RAID1 and configure the partitions:


Step 1 From the SUSE Suggested Partitioning page, select **Expert Partitioner > Next**.



Step 2 From the SUSE Add menu, select **Add > RAID**.

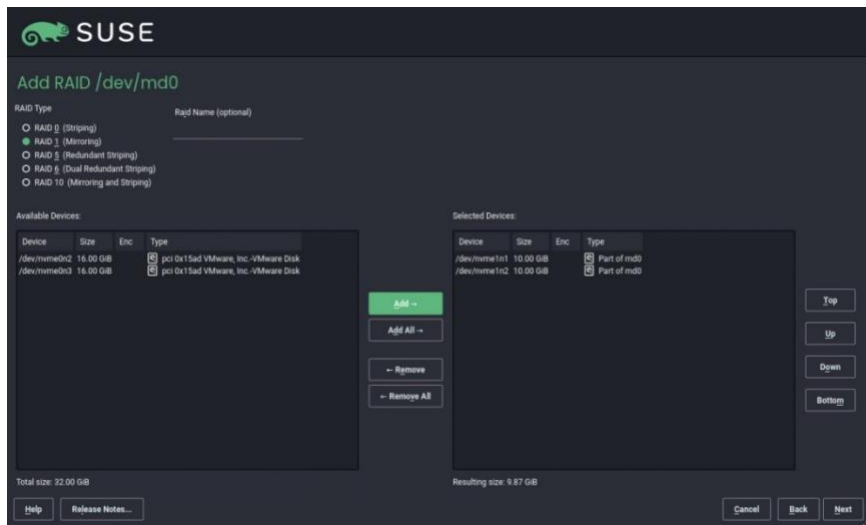


Step 3 From the SUSE Add RAID page, select **RAID 1 (Mirroring)** for the RAID Type.



Step 4 From the **Selected Devices** list, select two NVMe disks and click **Add**.

Step 5 Click **Next** to continue with the installation.



Manually Migrating the RAID Configuration Between Hosts

The following procedure describes how to migrate the RAID configuration manually between hosts.

Restoring a RAID Configuration from a Backup Configuration File

To restore a RAID configuration from a backup configuration file:

- Step 1** Periodically back up the configuration file `/etc/graid.conf` from the original host. Use `cp` or `scp` to move the configuration file to another system.
- Step 2** Set up the target host and ensure that the SupremeRAID™ service is stopped.

Note: If the target host already contains an installed and running SupremeRAID™ card, stop the service and copy the `graid.conf` file from the original system. On the original system, stop any running applications or unmount the mountpoint before starting the SupremeRAID™ service.

- Step 3** Move all the SSDs from the original host to the new host.

Step 4 Install the SupremeRAID™ driver on the new server. Stop the SupremeRAID™ service before copying the configuration backup file to the new host using the same path (/etc/graid.conf).

```
$ sudo systemctl stop graid
```

Step 5 Copy the configuration file

```
$ sudo cp graid.conf /etc/graid.conf
```

Step 6 If the original card also moved to the new host, start the SupremeRAID™ service directly.

```
$ sudo systemctl start graid
```

Otherwise, you must apply the new license if the card changed.

```
$ sudo graidctl apply license <LICENSE_KEY>
```

Restoring a RAID Configuration from SSD Metadata

The SupremeRAID™ system provides robust support for restoring RAID configurations from SSD metadata. This feature allows you to recover a RAID configuration quickly and easily in case of a failure or other issues. Perform the following procedure to restore the RAID configuration and get the SupremeRAID™ system back online.

To restore a RAID configuration from an SSD's metadata:

Step 1 Set up the target host and make sure that the SupremeRAID™ service is stopped.

Note: If the target host already contains an installed and running SupremeRAID™ card, stop the service the SupremeRAID™ service before restoring the configuration. On the original system, stop any running applications or unmount the mountpoint before starting the SupremeRAID™ service.

Step 2 Move all the SSDs from the original host to the new host.

Step 3 Install the SupremeRAID™ driver on the new server and stop the SupremeRAID™ service before restoring the configuration file.

```
$ sudo systemctl stop graid
```

Step 4 Run the restore command and restore the configuration file from SSD's metadata.

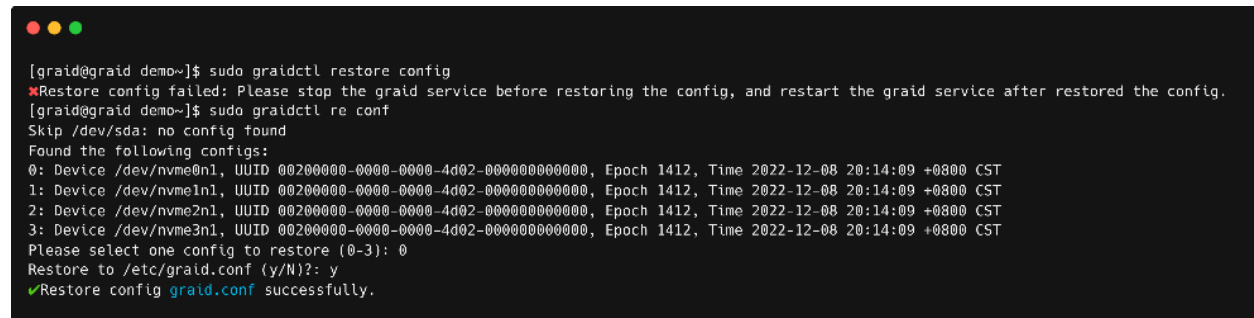
```
$ sudo graidctl restore config
```


Step 5 If the original card also moved to the new host, start the SupremeRAID™ service directly.

```
$ sudo systemctl start graid
```

Otherwise, you must apply the new license if the card changed.

```
$ sudo graidctl apply license <LICENSE_KEY>
```



```
[graid@graid demo~]$ sudo graidctl restore config
✖Restore config failed: Please stop the graid service before restoring the config, and restart the graid service after restored the config.
[graid@graid demo~]$ sudo graidctl re conf
Skip /dev/sda: no config found
Found the following configs:
0: Device /dev/nvme0n1, UUID 00200000-0000-0000-4d02-000000000000, Epoch 1412, Time 2022-12-08 20:14:09 +0800 CST
1: Device /dev/nvme1n1, UUID 00200000-0000-0000-4d02-000000000000, Epoch 1412, Time 2022-12-08 20:14:09 +0800 CST
2: Device /dev/nvme2n1, UUID 00200000-0000-0000-4d02-000000000000, Epoch 1412, Time 2022-12-08 20:14:09 +0800 CST
3: Device /dev/nvme3n1, UUID 00200000-0000-0000-4d02-000000000000, Epoch 1412, Time 2022-12-08 20:14:09 +0800 CST
Please select one config to restore (0-3): 0
Restore to /etc/graid.conf (y/N)? y
✔Restore config graid.conf successfully.
```

Restarting the SupremeRAID™ Service After Upgrading the System Kernel

If the SupremeRAID™ service does not start properly after upgrading the kernel, reinstall the SupremeRAID™ pre-installer and the installer to ensure that they are configured properly for the new kernel environment.

To reinstall the SupremeRAID™ pre-installer and installer on new kernel, follow these steps:

Step 1 Download the latest version of the SupremeRAID™ pre-installer and installer from the Graid Technology website.

Step 2 Open a terminal window and log in to the system as a user with root privileges.

Step 3 Use the **cd** command to navigate to the directory where the downloaded installer files are located.

Step 4 Run the graid-sr-pre-installer and follow the on-screen instructions to complete the pre-installation process.

Step 5 Run the graid-sr-installer and follow the on-screen instructions to complete the installation process.

After the SupremeRAID™ pre-installer and installer are installed successfully, use the **restart** command to restart the SupremeRAID™ service and confirm that it is running properly in the new kernel environment.

```
sudo systemctl restart graid
```

Obtaining SMART Information from Devices

Self-Monitoring, Analysis and Reporting Technology (SMART) data is a set of metrics and parameters that SSDs collect and monitor to assess their health and performance. Although the specific information included in the SMART data varies by manufacturer and drive model, it typically reports on the temperature, available spare capacity, power-on hours, error rates, and other details that are used to monitor the health of the SSD and predict its future performance.

By monitoring the SMART data for an SSD, you can identify a potential issue or degradation of the drive before it becomes a serious problem.

To check the SMART information for the gpd device using the NVMe `smart-log` or `smartctl` command, follow these steps:

Step 1 Open a terminal window and log in to the system with administrative privileges.

Step 2 Use the `list physical drives` command to list the available physical devices and identify the device name for the gpd device. For example, the device name might be `/dev/gpdx`.

To list the available physical drives and identify the device name:

```
$ sudo graidctl list physical_drive
```

Step 3 Use the `nvme` command to display the SMART data for the gpd device:

```
$ sudo nvme smart-log /dev/gpd<#>
```

Alternatively, you can use the `smartctl` command to display the SMART data for the gpd device:

```
$ sudo smartctl -d nvme -a /dev/gpd<#>
```

A detailed report of the SMART data for the gpd device, including the temperature, available spare capacity, and other details, appears. Use this information to monitor the health and performance of the device and to diagnose any potential issues.

Note: The specific steps and commands used to display SMART data may vary, depending on your system and the version of the `nvme` or `smartctl` command in use. Be sure to use the correct device name for the gpd device in the command.

The following figure shows an output example using nvme smart-log.

```
[graid@graid-demo ~]$ sudo graidctl list physical_drive
✔List physical drive successfully.
```

| PD ID (4) | DG ID | DEVICE PATH | NQN/WWID | MODEL | CAPACITY | SLOT ID | NUMA NODE | STATE |
|-----------|-------|-------------|-------------------------------|----------------|----------|---------|-----------|--------|
| 0 | 0 | /dev/gpd3 | nqn.2019-08.org.qemu:NVME0001 | QEMU NVMe Ctrl | 9.9 GiB | N/A | 0 | ONLINE |
| 1 | 1 | /dev/gpd0 | nqn.2019-08.org.qemu:NVME0002 | QEMU NVMe Ctrl | 9.9 GiB | N/A | 0 | ONLINE |
| 2 | 1 | /dev/gpd1 | nqn.2019-08.org.qemu:NVME0003 | QEMU NVMe Ctrl | 9.9 GiB | N/A | 0 | ONLINE |
| 3 | 1 | /dev/gpd2 | nqn.2019-08.org.qemu:NVME0004 | QEMU NVMe Ctrl | 9.9 GiB | N/A | 0 | ONLINE |

```
[graid@graid-demo ~]$ sudo nvme smart-log /dev/gpd0
Smart Log for NVME device:gpd0 namespace-id:ffffff
critical_warning      : 0
temperature           : 50 C
available_spare       : 0%
available_spare_threshold : 0%
percentage_used       : 0%
endurance_group      : 0
critical_warning_summary: 0
data_units_read      : 130566489
data_units_written   : 371
host_read_commands   : 510046973
host_write_commands  : 14028
controller_busy_time : 0
power_cycles          : 0
power_on_hours       : 68
unsafe_shutdowns     : 0
media_errors         : 0
num_err_log_entries  : 0
Warning Temperature Time : 0
Critical Composite Temperature Time : 0
Thermal Management T1 Trans Count : 0
Thermal Management T2 Trans Count : 0
Thermal Management T1 Total Time : 0
Thermal Management T2 Total Time : 0
```

The following figure shows an output example using smartctl.

```
[graid@graid-demo ~]$ sudo graidctl list physical_drive
✓List physical drive successfully.
```

| PD ID (4) | DG ID | DEVICE PATH | NON/WWID | MODEL | CAPACITY | SLOT ID | NUMA NODE | STATE |
|-----------|-------|-------------|-------------------------------|----------------|----------|---------|-----------|--------|
| 0 | 0 | /dev/gpd3 | nqn.2019-08.org.qemu:NVME0001 | QEMU NVMe Ctrl | 9.9 GiB | N/A | 0 | ONLINE |
| 1 | 1 | /dev/gpd0 | nqn.2019-08.org.qemu:NVME0002 | QEMU NVMe Ctrl | 9.9 GiB | N/A | 0 | ONLINE |
| 2 | 1 | /dev/gpd1 | nqn.2019-08.org.qemu:NVME0003 | QEMU NVMe Ctrl | 9.9 GiB | N/A | 0 | ONLINE |
| 3 | 1 | /dev/gpd2 | nqn.2019-08.org.qemu:NVME0004 | QEMU NVMe Ctrl | 9.9 GiB | N/A | 0 | ONLINE |

```
[graid@graid-demo ~]$ sudo smartctl -d nvme -a /dev/gpd0
smartctl 7.1 2020-04-05 r5049 [x86_64-linux-4.18.0-348.7.1.el8_5.x86_64] (local build)
Copyright (C) 2002-19, Bruce Allen, Christian Franke, www.smartmontools.org

=== START OF INFORMATION SECTION ===
Model Number:          QEMU NVMe Ctrl
Serial Number:         NVME0002
Firmware Version:      1.0
PCI Vendor ID:         0x1b36
PCI Vendor Subsystem ID: 0x1af4
IEEE OUI Identifier:   0x525400
Controller ID:         0
Number of Namespaces:  256
Local Time is:         Thu Dec  8 14:28:52 2022 CST
Firmware Updates (0x03): 1 Slot, Slot 1 R/0
Optional Admin Commands (0x000a): Format NS_Mngmt
Optional NVM Commands (0x015d):  Comp DS_Mngmt Wr_Zero Sav/Sel_Feat Timestmp *Other*
Maximum Data Transfer Size: 128 Pages
Warning Comp. Temp. Threshold: 70 Celsius
Critical Comp. Temp. Threshold: 100 Celsius

Supported Power States
St Op   Max Active   Idle   RL RT WL WT  Ent_Lat  Ex_Lat
0 +    25.00W    -      -    0 0 0 0     16      4

=== START OF SMART DATA SECTION ===
SMART overall-health self-assessment test result: PASSED

SMART/Health Information (NVMe Log 0x02)
Critical Warning:         0x00
Temperature:              50 Celsius
Available Spare:          0%
Available Spare Threshold: 0%
Percentage Used:          0%
Data Units Read:          130,566,489 [66.8 TB]
Data Units Written:        371 [189 MB]
Host Read Commands:       510,046,973
Host Write Commands:      14,028
Controller Busy Time:     0
Power Cycles:              0
Power On Hours:           68
Unsafe Shutdowns:         0
Media and Data Integrity Errors: 0
Error Information Log Entries: 0
Warning Comp. Temperature Time: 0
Critical Comp. Temperature Time: 0

Error Information (NVMe Log 0x01, max 1 entries)
No Errors Logged
```

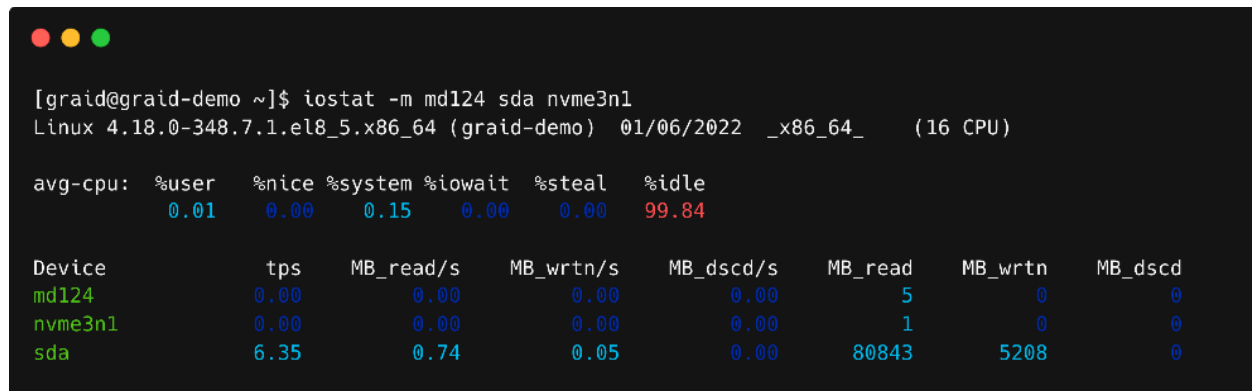
Monitoring System Input/Output Statistics for Devices Using iostat

The sysstat package contains the tools most commonly used to monitor I/O statistics in Linux systems. The sysstat package includes the iostat tool, which monitors system I/O device loading by observing the time the devices are active relative to their average transfer rates. The **iostat** command generates reports that allow you to fine-tune the system configuration to better balance the I/O load between physical disks.

For example, to monitor specific devices and display statistics in megabytes per second (Mbps), issue the following command:

```
$ iostat -m md124 sda nvme0n1
```

The following figure shows an output example.



```
[graid@graid-demo ~]$ iostat -m md124 sda nvme3n1
Linux 4.18.0-348.7.1.el8_5.x86_64 (graid-demo) 01/06/2022 _x86_64_ (16 CPU)

avg-cpu:  %user   %nice %system %iowait  %steal   %idle
           0.01    0.00    0.15    0.00    0.00   99.84

Device            tps    MB_read/s    MB_wrtn/s    MB_dscd/s    MB_read    MB_wrtn    MB_dscd
md124              0.00         0.00         0.00         0.00         5          0          0
nvme3n1            0.00         0.00         0.00         0.00         1          0          0
sda                6.35         0.74         0.05         0.00       80843       5208         0
```

sysstat Versions v12.3.3 and Later

For sysstat versions v12.3.3 and later, the iostat tool includes an alternative directory feature that allows you to specify the directory from which to read device statistics.

- Add a **+f** parameter to the tool and use the `/sys/devices/virtual/graid/graid` sysfs device path to read device statistics from both the standard kernel files and the files in the alternative directory.
- Add a **-f** parameter to the tool and use the `/sys/devices/virtual/graid/graid` sysfs device path to read device statistics from the files in the alternative directory.

The following figure shows an alternative directory description from the iostat manual page.

```

-f directory
+f directory
Specify an alternative directory for iotop to read device statistics. Option -f tells iotop to use only the files located in the alternative directory, whereas option +f tells it to use both the standard kernel files and the files located in the alternative directory to read device statistics.

directory is a directory containing files with statistics for devices managed in userspace. It may contain:
- a "diskstats" file whose format is compliant with that located in /proc,
- statistics for individual devices contained in files whose format is compliant with that of files located in /sys.

In particular, the following files located in directory may be used by iotop:

directory/block/device/stat
directory/block/device/partition/stat

partition files must have an entry in directory/dev/block/ directory, e.g.:

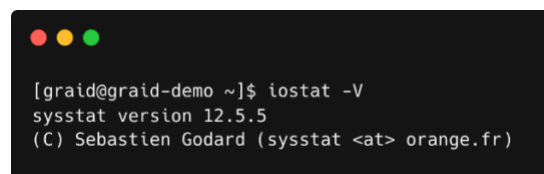
directory/dev/block/major:minor --> ../block/device/partition

```

To check the `iotop` version, issue the following command:

```
$ iostat -V
```

The following figure shows an output example.



```

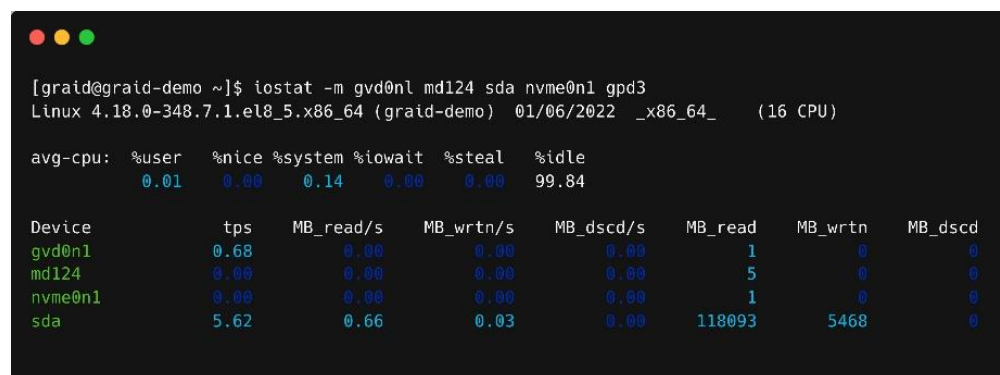
[graid@graid-demo ~]$ iostat -V
sysstat version 12.5.5
(C) Sebastien Godard (sysstat <at> orange.fr)

```

The `gpd#` statistics are not displayed in the `iotop` report without appending the `+f` parameter and defining the `sysfs` path.

```
$ iostat -m +f /sys/devices/virtual/graid/graid gdg0n1 md124 sda nvme0n1 gpd3
```

The following figure shows an output example.



```

[graid@graid-demo ~]$ iostat -m gvd0n1 md124 sda nvme0n1 gpd3
Linux 4.18.0-348.7.1.el8_5.x86_64 (graid-demo) 01/06/2022 _x86_64_ (16 CPU)

avg-cpu:  %user   %nice %system %iowait  %steal   %idle
           0.01    0.00    0.14    0.00    0.00   99.84

Device            tps    MB_read/s    MB_wrtn/s    MB_dscd/s    MB_read    MB_wrtn    MB_dscd
gvd0n1             0.68         0.00         0.00         0.00         1           0           0
md124              0.00         0.00         0.00         0.00         5           0           0
nvme0n1            0.00         0.00         0.00         0.00         1           0           0
sda                5.62         0.66         0.03         0.00       118093       5468           0

```

The `gpd#` statistics are displayed when the `+f` parameter is appended and the `sysfs` path is defined.

```
$ iostat -m +f /sys/devices/virtual/graid/graid gdg0n1 md124 sda nvme0n1 gpd3
```

The following figure shows an output example.

```
[graid@graid-demo ~]$ iostat -m +f /sys/devices/virtual/raid/graid/graid gvd0n1 md124 sda nvme0n1 gpd3
Linux 4.18.0-348.7.1.el8_5.x86_64 (graid-demo) 01/06/2022 _x86_64_ (16 CPU)

avg-cpu:  %user   %nice %system %iowait  %steal   %idle
           0.01    0.00    0.15    0.00    0.00   99.84

Device            tps    MB_read/s    MB_wrtn/s    MB_dscd/s    MB_read    MB_wrtn    MB_dscd
gpd3              0.00         0.00         0.00         0.00         9          0          0
gvd0n1            0.00         0.00         0.00         0.00         2          0          0
md124             0.00         0.00         0.00         0.00         5          0          0
nvme0n1          0.00         0.00         0.00         0.00         1          0          0
sda              6.22         0.72         0.05         0.00       80853       5208         0
```

sysstat Versions Prior to v12.3.3

For operating systems with sysstat versions prior to v12.3.3 (for example, CentOS), Graid Technology provides an alternate tool called giostat to display device statistics.

In the following example, the operating system version of iostat is prior to v12.3.3.

```
$ sudo yum list --installed |grep sysstat
```

The following figure shows an output example.

```
[graid@graid-demo ~]$ sudo yum list --installed |grep sysstat
sysstat.x86_64                               11.7.3-6.el8                                @appstream
```

The giostat and iostat tools are very similar and their usage is the same. Set the parameter preferences using giostat. The following figure shows an output example.

```
[graid@graid-demo ~]$ sudo graidctl list physical drive;sudo graidctl list drive group ;sudo graidctl list virtual drive
✓List physical drive successfully.
```

| PD ID (5) | DG ID | DEVICE PATH | QN/QWID | MODEL | CAPACITY | SLOT ID | STATE |
|-----------|-------|-------------|--|--------------|----------|---------|-------------------|
| 0 | 0 | /dev/gpd0 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z0G0A001T1L8 | KCM61VUL3T20 | 3.2 TB | 12 | ONLINE |
| 1 | 0 | /dev/gpd3 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z010A004T1L8 | KCM61VUL3T20 | 3.2 TB | 19 | ONLINE |
| 2 | 0 | /dev/gpd2 | nqn.2019-10.com.kioxia:KCM61VUL3T20:X0X0A01ET1L8 | KCM61VUL3T20 | 3.2 TB | 18 | ONLINE |
| 3 | 0 | /dev/gpd1 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z030A04HT1L8 | KCM61VUL3T20 | 3.2 TB | 8 | ONLINE |
| 4 | N/A | /dev/gpd4 | nqn.2019-10.com.kioxia:KCM61VUL3T20:Z030A038T1L8 | KCM61VUL3T20 | 3.2 TB | 0 | UNCONFIGURED_GOOD |

```
✓List drive group successfully.
```

| DG ID | MODE | VD NUM | CAPACITY | FREE | USED | STATE |
|-------|-------|--------|----------|--------|-------|---------|
| 0 | RAID6 | 4 | 6.4 TB | 6.4 TB | 25 GB | OPTIMAL |

```
✓List virtual drive successfully.
```

| VD ID (4) | DG ID | SIZE | DEVICE PATH | STATE | EXPORTED |
|-----------|-------|--------|-------------|--------|----------|
| 0 | 0 | 10 GB | /dev/gvd0n1 | RESYNC | No |
| 1 | 0 | 5.0 GB | /dev/gvd1n1 | RESYNC | No |
| 2 | 0 | 5.0 GB | /dev/gvd2n1 | RESYNC | No |
| 3 | 0 | 5.0 GB | /dev/gvd3n1 | RESYNC | No |

```
[graid@graid-demo ~]$ giostat -m gvd0n1 gpd3 nvme10n1 sda
Linux 4.18.0-348.2.1.el8_5.x86_64 (graid-demo) 01/06/2022 _x86_64_ (128 CPU)
```

```
avg-cpu:  %user   %nice %system %iowait  %steal   %idle
           0.02    0.00    0.04    0.00    0.00   99.93
```

| Device | tps | MB read/s | MB wrtn/s | MB dscd/s | MB read | MB wrtn | MB dscd |
|----------|---------|-----------|-----------|-----------|---------|---------|---------|
| gpd3 | 1449.98 | 3.79 | 4.19 | 0.00 | 3355542 | 3707736 | 0 |
| gvd0n1 | 0.05 | 0.01 | 0.00 | 0.00 | 9530 | 0 | 0 |
| nvme10n1 | 0.00 | 0.00 | 0.00 | 0.00 | 1 | 0 | 0 |
| sda | 0.00 | 0.00 | 0.00 | 0.00 | 15 | 0 | 0 |

Setting Up the Auto-mount File Systems on Linux Using the SupremeRAID™ Driver

To set up the auto-mount file systems on Linux using the SupremeRAID™ driver:

Step 1 Create a virtual drive.

```
$ sudo graidctl create virtual_drive <DG_ID> [size] [flags]
```

Step 2 Format the virtual drive and create a mount point for it.

```
$ sudo mkdir /mnt/<name-of-the-drive>
$ sudo mkfs.<file-system-type> /dev/gdgXnY
$ sudo mount /dev/gdgXnY /mnt/<name-of-the-drive>/
```


Step 3 Obtain the name, UUID, and file system type.

```
$ ls -l /dev/disk/by-id/
```

Step 4 Edit the /etc/fstab file:

A Edit the /etc/fstab file.

```
$ sudo vim /etc/fstab
```

B Append one line of code to the end of the file using the following format:

```
UUID=<uuid-of-the-drive> <mount-point> <file-system-type> x-
systemd.requires=graid.service,nofail <dump> <pass>
```

The following figure shows an output example.

```
[root@graid-demo ~]# ls -l /dev/disk/by-id/
total 0
lrwxrwxrwx. 1 root root 12 Sep  8 06:27 gdg-eui.00abcdef00136d5b65ald3d7ecb5b8ad -> ../../gdg0n1
lrwxrwxrwx. 1 root root 12 Sep  8 06:27 gdg-GRAID-SR_96BCDBC839F109EE_1 -> ../../gdg0n1
lrwxrwxrwx. 1 root root 10 Sep  6 05:09 lvm-pv-uuid-cjIZ8z-5SmL-8NmF-z6lA-1zlk-J5DT-HGFlnS -> ../../sda3
lrwxrwxrwx. 1 root root  9 Sep  7 23:12 md-name-graid-demo:0 -> ../../md0
lrwxrwxrwx. 1 root root  9 Sep  7 23:12 md-uuid-636e39c5:cbfa794e:91f4dd06:e8fbc6be -> ../../md0
lrwxrwxrwx. 1 root root 13 Sep  7 23:12 nvme-
nvme.lb36-4e564d4530303032-51454d55204e564d65204374726c-00000001 -> ../../nvme0n1
lrwxrwxrwx. 1 root root 13 Sep  7 23:12 nvme-
nvme.lb36-4e564d4530303034-51454d55204e564d65204374726c-00000001 -> ../../nvme1n1
lrwxrwxrwx. 1 root root 13 Sep  7 23:12 nvme-QEMU_NVMe_Ctrl_NVME0002 -> ../../nvme0n1
lrwxrwxrwx. 1 root root 13 Sep  7 23:12 nvme-QEMU_NVMe_Ctrl_NVME0004 -> ../../nvme1n1

[root@graid-demo ~]# sudo vim /etc/fstab
#
# /etc/fstab
# Created by anaconda on Thu May 18 23:02:31 2023
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
/dev/mapper/rhel-root / xfs defaults 0 0
UUID=f6f00b7c-87d8-472a-90d1-41b73372b792 /boot xfs defaults 0 0
UUID=6C6D-B3E9 /boot/efi vfat umask=0077,shortname=winnt 0 0
/dev/mapper/rhel-swap swap swap defaults 0 0

/dev/disk/by-id/gdg-GRAID-SR_96BCDBC839F109EE_1 /mnt/graid_demo ext4 x-
systemd.requires=graid.service,nofail 0 0

#UUID=9c2ca3e2-6adc-44cc-926a-4125282cef15 /mnt/graid_demo1.5 xfs x-systemd.requires=graid.service,nofail
0 0
```

Step 5 Remove the device line and reboot the system.

```
$ sudo vim /etc/fstab
```

The following figure shows an output example.

```
[root@graid-demo ~]# ls -l /dev/disk/by-id/
total 0
lrwxrwxrwx. 1 root root 12 Sep  8 06:27 gdg-eui.00abcdef00136d5b65a1d3d7ecb5b8ad -> ../../gdg0n1
lrwxrwxrwx. 1 root root 12 Sep  8 06:27 gdg-GRAID-SR_96BCDBC839F109EE_1 -> ../../gdg0n1
lrwxrwxrwx. 1 root root 10 Sep  6 05:09 lvm-pv-uuid-cjIZ8z-5SmL-8NmF-z6lA-1z1k-J5DT-HGFlnS -> ../../sda3
lrwxrwxrwx. 1 root root  9 Sep  7 23:12 md-name-graid-demo:0 -> ../../md0
lrwxrwxrwx. 1 root root  9 Sep  7 23:12 md-uuid-636e39c5:cbfa794e:91f4dd06:e8fbc6be -> ../../md0
lrwxrwxrwx. 1 root root 13 Sep  7 23:12 nvme-
nvme.1b36-4e564d4530303032-51454d55204e564d65204374726c-00000001 -> ../../nvme0n1
lrwxrwxrwx. 1 root root 13 Sep  7 23:12 nvme-
nvme.1b36-4e564d4530303034-51454d55204e564d65204374726c-00000001 -> ../../nvme1n1
lrwxrwxrwx. 1 root root 13 Sep  7 23:12 nvme-QEMU_NVMe_Ctrl_NVME0002 -> ../../nvme0n1
lrwxrwxrwx. 1 root root 13 Sep  7 23:12 nvme-QEMU_NVMe_Ctrl_NVME0004 -> ../../nvme1n1

[root@graid-demo ~]# sudo vim /etc/fstab
#
# /etc/fstab
# Created by anaconda on Thu May 18 23:02:31 2023
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
/dev/mapper/rhel-root / xfs defaults 0 0
UUID=f6f00b7c-87d8-472a-90d1-41b73372b792 /boot xfs defaults 0 0
UUID=6C6D-B3E9 /boot/efi vfat umask=0077,shortname=winnt 0 0
/dev/mapper/rhel-swap swap swap defaults 0 0

#/dev/disk/by-id/gdg-GRAID-SR_96BCDBC839F109EE_1 /mnt/graid_demo ext4 x-
systemd.requires=graid.service,nofail 0 0

#UUID=9c2ca3e2-6adc-44cc-926a-4125282cef15 /mnt/graid_demo1.5 xfs x-systemd.requires=graid.service,nofail
0 0
~
```

Note: To disable the automount point or delete the virtual drive, edit the `/etc/fstab` file to delete/comment that entry, and then reboot the system.

ESXi Virtual Machine Support Using GPU Passthrough

You can create virtual machines with SupremeRAID™ support to maximize performance.

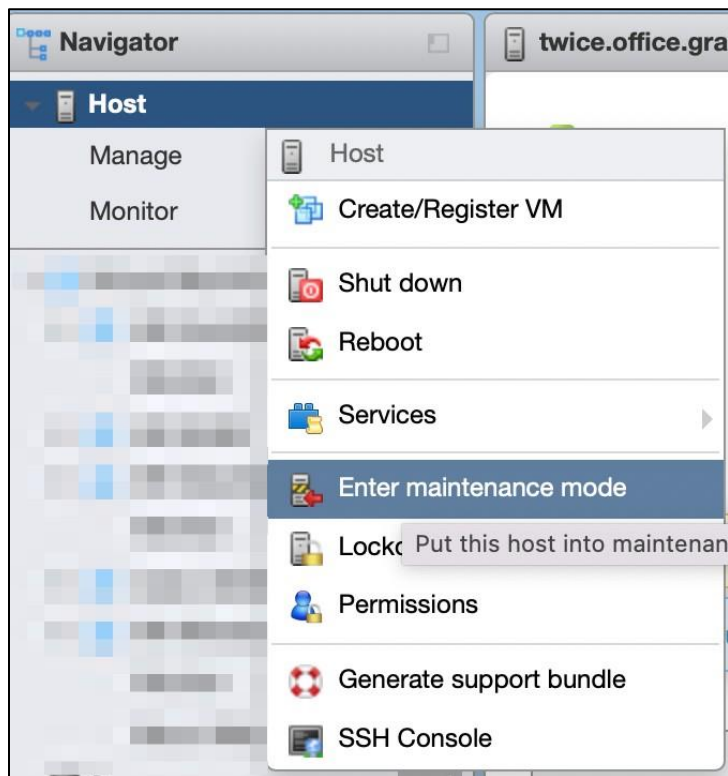
The following procedure describes how to set a single VM with SupremeRAID™. This setup is for use only within a single virtual machine and cannot be shared from the volume back to ESXi to a datastore for other virtual machines.

Hypervisor VMware support is ESXi 7.0U3.

Configuring Hosts for NVIDIA GPU Device Passthrough

Setting the ESXi Host in Maintenance Mode

From the Navigator menu, select **Host > Enter maintenance mode**.



Managing PCI Device Passthrough

- Step 1** From the **Navigator** menu, select **Manage > Hardware > PCI Devices**. The Passthrough Configuration page appears, listing all available passthrough devices.
- Step 2** Select the NVIDIA T1000 (Quadro T1000 Mobile) and its Audio device.
- Step 3** Click Toggle passthrough.
- Step 4** Confirm that the Passthrough status is **Active**.

| Address | Description | SR-IOV | Passthrough | Hardware Label |
|--|--|-------------|-------------|----------------|
| 0000:40:03.1 | Advanced Micro Devices, Inc. [AMD] Starship/Matisse GPP Bridge | Not capable | Not capable | |
| <input checked="" type="checkbox"/> 0000:42:00.1 | nVidia Corporation Audio device | Not capable | Active | |
| <input checked="" type="checkbox"/> 0000:42:00.0 | NVIDIA Corporation TU117GLM [Quadro T1000 Mobile] | Not capable | Active | |
| 0000:40:04.0 | Advanced Micro Devices, Inc. [AMD] Starship/Matisse PCIe Dummy Host Bridge | Not capable | Not capable | |
| 0000:40:05.0 | Advanced Micro Devices, Inc. [AMD] Starship/Matisse PCIe Dummy Host Bridge | Not capable | Not capable | |
| 0000:40:07.0 | Advanced Micro Devices, Inc. [AMD] Starship/Matisse PCIe Dummy Host Bridge | Not capable | Not capable | |
| 0000:40:07.1 | Advanced Micro Devices, Inc. [AMD] Starship/Matisse Internal PCIe GPP Bridg... | Not capable | Not capable | |

Note: If you move the SupremeRAID™ card to a different hardware slot or plan to do so, you **MUST** cancel its passthrough before shutting down the ESXi server. After the hardware change, you **MUST** set up the passthrough again; otherwise, the virtual machine will not recognize the PCIe device properly.

Configuring Virtual Machines

Attaching PCI Devices to the Virtual Machine

To attach PCI devices to the virtual machine:


- Step 1** From the Edit VM setting page, select **Virtual Hardware > Add other device > PCI device**.
- Step 2** Select Quadro T1000 and its Audio device as the two PCI devices.

Note: When the T1000 PCI device is assigned to the virtual machine, you must set the memory reservation to accommodate the fully configured memory size.



Step 3 Select Virtual Hardware > Memory.


Step 4 Check Reserve all guest memory (All locked).



Enabling Point-to-Point (P2P) on the Virtual Machine

Enabling P2P on the virtual machine optimizes performance. To enable P2P on the virtual machine:

Step 1 From the Edit VM setting page, select VM Options > Advanced > Configuration Parameters > Edit Configuration....




Step 2 Add the following two parameters:

```
hypervisor.cpuid.v0 = "FALSE"
pciPassthru.allowP2P = "TRUE" pciPassthru.use64bitMMIO= "TRUE"
```

Step 3 From the Edit VM setting page, select VM Options > Boot Options > Firmware > EFI.

Step 4 Uncheck Whether or not to enable UEFI secure boot for this VM.



Using Self-Encrypting Drives

A self-encrypting drive (SED) uses native full-disk encryption. SupremeRAID™ supports SEDs and SED key management. When the SED key is configured, SupremeRAID™ uses the imported key to unlock the SED.

Before configuring a SED, observe the following guidelines:

- Configure the SED key using the `graidctl` tool before creating the physical drives.
- Only NVMe devices are supported.
- Only the global range parameter is supported.

Importing a Single SED Key Using NQN/WWID

To import a single SED key using NQN/WWID, issue the following command:

```
$ sudo graidctl edit config sed <NQN/WWID>
```

The following figure shows an example.

```
graid@graid:~$ sudo RAIDCTL edit config sed nqn.2014-08.org.nvmexpress:uuid:52bbdb40-c5bf-f92d-9961-a6368e845bfd
Enter Key:
✓ Edit config sed successfully.
```

Importing a Batched SED Key Using NQN/WWID

To import a batched SED key using NQN/WWID, issue the following command:

```
$ sudo RAIDCTL edit config sed file <filename> file content format:
<NQN1/WWID1>, <KEY1>
<NQN1/WWID1>, <KEY2>
...
<NQNn/WWIDn>, <KEYn>
```

Displaying SED Key Information

To display SED key information, issue the following command:

```
$ sudo RAIDCTL describe config sed
```

The following figure shows an example.

```
graid@graid:~$ sudo RAIDCTL describe config sed
Totally 1 GUIDs have SED key:
nqn.2014-08.org.nvmexpress:uuid:52bbdb40-c5bf-f92d-9961-a6368e845bfd
```

Deleting SED Keys

To delete a SED key, issue the following command:

```
$ RAIDCTL delete config sed <GUID>
```

The following figure shows an example.

```
graid@graid:~$ sudo graidctl delete config sed nqn.2014-08.org.nvmexpress:uuid:52bbdb40-c5bf-f92d-9961-a6368e845bfd
✔ Delete config sed successfully.
```

To delete all SED keys, issue the following command:

```
$ graidctl delete config sed all
```

The following figure shows an example.

```
graid@graid:~$ sudo graidctl delete config sed all
Do you really want to delete all SED key?
Repeat IMEANTODELETEALL to continue: IMEANTODELETEALL
✔ Delete config sed successfully.
```


TROUBLESHOOTING

Sequential Read Performance is Not as Expected on a New Drive Group

Unlike SAS/SATA hard drives, many NVMe SSDs support the de-allocate dataset management command. Using this command, you can reset all data in the NVMe SSD immediately, eliminating the need to synchronize data between physical drives when creating a drive group.

For other SSDs, however, the performance is not as expected when reading unwritten sectors after issuing the de-allocate dataset management command. While this behavior also impacts the performance of the new drive group, it does not affect the applications because they do not read sectors that do not contain data.

To test SupremeRAID™ performance, write the entire virtual drive sequentially using a large block size.

Kernel Log Message "failed to set APST feature (-19)" Appears When Creating Physical Drives

Some NVMe SSD models might display a "failed to set APST feature (-19)" message in the kernel log when creating the physical drive.

When SupremeRAID™ creates the physical drive, the SSD is unbound from the operating system so the SupremeRAID™ can control the SSD. When the APST feature is enabled during the unbinding process, the NVMe driver tries and fails to set the APST state to SSD and the error message is issued. This message is expected and can be ignored. SupremeRAID™ is working normally.

Decoding LED Patterns on the Backplane

You might notice that the HDD/SSD activity indicator blink pattern is different on SupremeRAID™ than on traditional RAID cards.

SupremeRAID™ does not require a buffering or caching mechanism to improve read/write performance as do traditional RAID cards. This feature causes SupremeRAID™ indicators to blink differently than traditional RAID cards.

Received "The arch of the controller and graid software mismatched" Message When Applying License

To activate the SupremeRAID™ server with your license key, it's essential to install the correct driver version that matches your specific SupremeRAID™ model. If the incorrect version is installed, the following error message appears when you try to activate the SupremeRAID™ server with a license key: **Apply license failed: The arch of the controller and graid software mismatched.**

To ascertain which model you installed, use the command **graidctl version**. Issuing this command displays the model information at the end of the string.

```
001 -> SupremeRAID™ SR-1001
000 -> SupremeRAID™ SR-1000
010 -> SupremeRAID™ SR-1010
```

The following figure shows an example of the message.

```
root@graid:/home/graid# graidctl version
✓Graidctl version successfully.
graidctl version:          1.5.0-rc1-644.ga82e0d9a.010
graid_server version:     1.5.0-rc1-644.ga82e0d9a.010
root@graid:/home/graid# graidctl apply license XXXXXXXXXXXXXXXX-XXXXXXXXXXXXXXXXXX
✖Apply license failed: No controller found for this license key
```

If you receive this message, uninstall the incorrect driver, and then install the correct one.

Step 1 Stop graid service.

```
$ sudo systemctl stop graid
```

Step 2 Unload the kernel model of graid.

```
$ sudo rmmmod graid_nvidia graid
```

Step 3 Uninstall the package using the command appropriate for your operating system:

- For Centos, Rocky Linux, AlmaLinux, RHEL, openSUSE, and SLES:

```
$ sudo rpm -e graid-sr
```

- For Ubuntu:

```
$ sudo dpkg -r graid-sr
```

Step 4 Confirm that the SupremeRAID™ module is unloaded.

```
$ sudo lsmod | grep graid
```

There should not be any output.

Step 5 Confirm that the SupremeRAID™ package is uninstalled using the command appropriate for your operating system:

- For Centos, Rocky Linux, AlmaLinux, RHEL, openSUSE, and SLES:

```
$ sudo rpm -qa | grep graid
```

- For Ubuntu:

```
$ sudo dpkg -l | grep graid
```

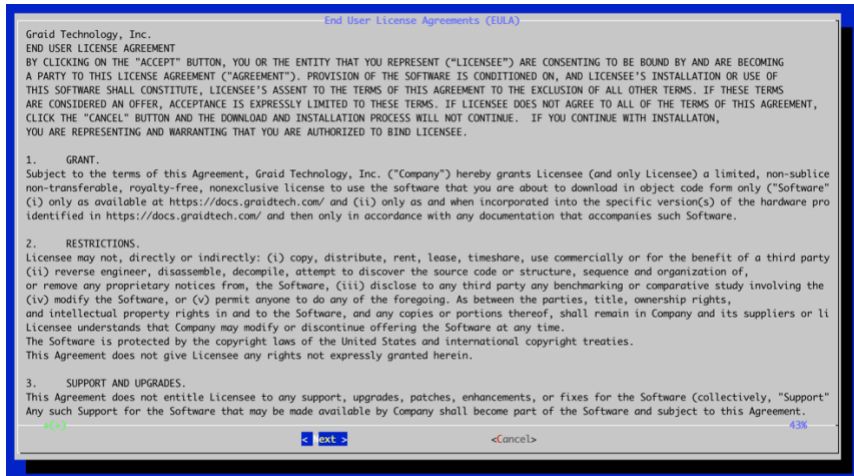
There should not be any output.

Step 6 Install the correct graid driver:

A At the Welcome page, select **Next** and click **Enter** to view the end-user license agreement.




B In the end-user license agreement, use the spacebar to scroll through the content. When you complete your review, select **Next** and click **Enter** to proceed.



C Type **accept**, click tab, select **Next**, and click **Enter** to accept the license agreement.



D Check the package version and click **NEXT**.



E To activate the software, apply the SupremeRAID™ license key.

```
$ sudo graidctl apply license <LICENSE_KEY>
```

SAFETY INFORMATION

English Version

CE Directives Declaration: NVIDIA Corporation hereby declares that this device complies with all material requirements and other relevant provisions of the 2014/30/EU and 2011/65/EU. A copy of the Declaration of Conformity may be obtained directly from NVIDIA GmbH (Bavaria Towers - Blue Tower, Einsteinstrasse 172, D-81677 Munich, Germany)

NVIDIA products are designed to operate safely when installed and used according to the product instructions and general safety practices. The guidelines included in this document explain the potential risks associated with equipment operation and provide important safety practices designed to minimize these risks. By carefully following the information contained in this document, you can protect yourself from hazards and create a safer environment.

This product is designed and tested to meet IEC 60950-1 and IEC 62368-1 Safety Standards for Information Technology Equipment. This also covers the national implementations of IEC 60950-1/62368-1 based safety standards around the world e.g. UL 62368-1. These standards reduce the risk of injury from the following hazards:

- Electric shock: Hazardous voltage levels contained in parts of the product
- Fire: Overload, temperature, material flammability
- Energy: Circuits with high energy levels (240-volt amperes) or potential as burn hazards.
- Heat: Accessible parts of the product at high temperatures.
- Chemical: Chemical fumes and vapors
- Radiation: Noise, ionizing, laser, ultrasonic waves

This device complies with part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) This device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

This product, as well as its related consumables and spares, complies with the reduction in hazardous substances provisions of the "India E-waste (Management and Handling) Rule 2016". It does not contain lead, mercury, hexavalent chromium, polybrominated biphenyls or polybrominated diphenyl ethers in concentrations exceeding 0.1 weight % and 0.01 weight % for cadmium, except for where allowed pursuant to the exemptions set in Schedule 2 of the Rule.

Retain and follow all product safety and operating instructions.

Always refer to the documentation supplied with your equipment. Observe all warnings on the product and in the operating instructions found on the product's User Guide.



This is a recycling symbol indicating that the product/battery cannot be disposed of in the trash and must be recycled according to the regulations and/or ordinances of the local community.



Hot surface warning. Contact may cause burns. Allow to cool before servicing.

Chinese Version

NVIDIA 产品在设计时充分考虑到操作安全性，可根据产品说明和常规安全做法进行安全安装和使用。本文档中包含的准则解释了设备操作所涉及的风险，并提供了最大限度降低这些风险的重要安全做法。请仔细阅读本文档中的信息并按要求操作，这样可保护您免遭受为显并创建一个更加安全的环境。

本产品按照信息技术设备安全标准 IEC 60950-1 和 IEC 62368-1 进行设计，并且经测试表明符合这些设备。此处所述标准也包括全球各国/地区实施的基于 IEC 60950-1/62368-1 的安全标准，例如 UL 62328-1。这些标准降低了因以下危险而受伤的风险：

- 电击：部分产品中包含的危险电压水平起火：超载、高温、可燃性材料
- 机械：锋利的边缘、活动部件、不稳定性
- 电源：高电压电路（240 伏安）或潜在的烧伤风险
- 高温：产品的可触及部分存在高温化学：化学烟雾和蒸气
- 辐射：噪音、电离、激光、超声波

请牢记并遵守所有产品安全和操作说明。请务必参考您的设备随附的说明文档。请注意产品上以及产品用户指南的操作说明中列

示的所有警告。



这是一个通用的回收标志，表示产品/电池不能以丢弃的方式处置，必须按造本地社区的法规和/或条例回收。



警告！表面高溫。接觸可能導致灼傷。請再冷卻後再使用。



产品中有害物质的名称及含量根据中国《电器电子产品有害物质限制使用管理办法》

| | | | | | | |
|--|---|---|---|---|---|---|
| 内存 | 0 | 0 | 0 | 0 | 0 | 0 |
| 结构回以及风扇 | x | 0 | 0 | 0 | 0 | 0 |
| 线材/连接器 | x | 0 | 0 | 0 | 0 | 0 |
| 焊接金属 | 0 | 0 | 0 | 0 | 0 | 0 |
| 助焊剂，锡膏，标签及耗材 | 0 | 0 | 0 | 0 | 0 | 0 |
| <p>本表格依据SJ/T 11364-2014的规定编制</p> <p>O：表示该有害物质在该部件所有的均质材料中的含量均在GB/T 26572 标准规定的限量要求以下。</p> <p>X：表示该有害物质至少在该部件的某一均质材料中的含量超出GB/T 26572标准规定的限量要求。</p> <p>此表中所有名称中含“X” 的部件均符合RoHS立法。</p> <p>注：环保使用期限的参考标识取决于产品正常工作的温度和湿度等条件</p> | | | | | | |

Chinese Version (TC)

在遵照產品說明與一般安全做法進行安裝與使用產品的情況下，NVIDIA 產品可安全地操作。本文件所列的準則說明與設備操作相關的潛在風險，同時也提供將這些風險降到最低的重要安全做法。謹慎遵守本文件中的資訊，您就可以避免危險並創造更安全的環境。

此產品係根據 Safety Standards for Information Technology Equipment(資訊技術設備安全標準) IEC 60950-1 和 IEC 62368-1 進

行設計與測試。同時也涵蓋全世界國家以 IEC 60950-1/62368-1 為根據的安全標準，例如 UL 62368-1。這些標準可降低下列危險造成的傷害的風險：

- 觸電危險：本產品部分零件的電壓等級具危險性
- 火災危險：超載、溫度、材料可燃性
- 機械危險：尖銳邊緣、移動零件、不穩定性

- 電燒力危險：電路電壓高（240 電壓）或具有潛在起火燃燒熱能危險：產品表面可能達到高溫，注意燙傷危機
- 化學危險：化學異味氣體與蒸氣
- 輻射危險：噪音、游離輻射、雷射、超音波

請保留並遵守所有產品安全與操作說明的相關規定。請務必參閱設備隨附的文件。請遵守產品上，和產品使用者只能中操作說明裡的警告規定。



此國際回收標誌表示此產品/電池不能棄置於垃圾桶中，必須根據當地社區的規範和/或法令回收。



表面高溫警告。接觸時可能燙傷。使用前請先降溫。

| 限用物質含有情況標示聲明書 | | | | | | |
|---------------|------------|---|---|-----|------|-------|
| 設備名稱：繪圖卡 | | | | | | |
| 單元 | 限用物質及其化學符號 | | | | | |
| | 鉛 | 汞 | 鎘 | 六價鉻 | 多溴聯苯 | 多溴二苯醚 |
| PCB板 | 0 | 0 | 0 | 0 | 0 | 0 |
| 結構閥以及風扇 | - | 0 | 0 | 0 | 0 | 0 |
| 連結器 | - | 0 | 0 | 0 | 0 | 0 |
| 被動電子零件 | - | 0 | 0 | 0 | 0 | 0 |
| 主動電子零件 | - | 0 | 0 | 0 | 0 | 0 |
| 內存 | 0 | 0 | 0 | 0 | 0 | 0 |
| 線材 | 0 | 0 | 0 | 0 | 0 | 0 |
| 焊接金屬 | 0 | 0 | 0 | 0 | 0 | 0 |
| 助焊劑、錫膏、標籤及耗材 | 0 | 0 | 0 | 0 | 0 | 0 |

備考1：0：係指該限用物質未超出百分比含量基準值

備考2：-：係指該限用物質為排除項目。

此表中所有名稱含“-”的部件均符合歐盟RoHS立法。

注：環保使用期限的參考標識取決於產品正常工作的溫度和濕度等條件

