



GRAID SupremeRAID™

User Guide for Windows

Version: 1.0.9

Copyright

Copyright © 2021–2022 GRAID Technology Inc. All Rights Reserved.

The term GRAID refers to GRAID Technology Inc. For more information, please visit www.graidtech.com.

GRAID reserves the right to make changes without further notice to any products or content herein to improve reliability, function, or design. GRAID 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’s or any third party’s intellectual property rights are conveyed hereunder.

Publication: July 19, 2022

Table of Contents

Table of Contents	1
Introduction	3
<i>GRAID SupremeRAID™ Specifications</i>	3
<i>Supported Operating Systems</i>	3
Installation	4
<i>Prerequisites</i>	4
<i>Installing on Windows</i>	4
<i>Upgrading the Software</i>	10
Management	11
<i>Overview of the GRAID SupremeRAID™ Software Module</i>	
RAID Components	11
Physical Drive (PD)	11
Drive Group (DG)	12
Virtual Drive (VD)	12
Overview of graidctl	13
Syntax	13
Managing Licenses	14
Applying the License	14
Checking License Information	15
Viewing Host Drive Information	16
Listing the NVMe Drives	16
Listing SAS/SATA Drives	17
Managing Physical Drives	18
Creating a Physical Drive	18
Listing the Physical Drives	19
Deleting Physical Drives	20
Describing a Physical Drive	20
Locating a Physical Drive	21
Marking a Physical Drive Online or Offline	21
Assigning a Hot Spare Drive	21
Replacing a Nearly Worn-Out or Broken SSD	21
Managing Drive Groups	23
Creating a Drive Group	23
Listing Drive Groups	24

Deleting Drive Groups	25
Describing a Drive Group	25
Setting the Drive Group Rebuild Speed	25
Locating the Physical Drives in the Drive Group	26
Degradation and Recovery	26
Rescue Mode	26
Managing Virtual Drives	27
Creating a Virtual Drive	27
Listing Virtual Drives	27
Deleting Virtual Drives	28
Create a RAID-5 Virtual Drive with 5 NVMe SSDs	28
Basic Troubleshooting	30
<i>Sequential Read Performance is not as Expected on a New Drive Group</i>	30
<i>Kernel Log Message "failed to set APST feature (-19)" Appears When Creating Physical Drives</i>	31
<i>Installer Error Message: "... nvml.dll not found..." Appears When Installing the SupremeRAID™ Driver</i>	
<i>Error Message: "GRAID service is not running"</i>	32
<i>Different LED Blink Patterns on the Backplane</i>	39
Appendix	40
<i>Manually Migrating the RAID Configuration Between Hosts</i>	40
<i>Enabling Virtual Machines with GPU Passthrough</i>	41
Configuring Hosts for NVIDIA GPU Device Passthrough	41
Configuring Virtual Machines	42
Safety Information	45
English version	45
Chinese version(SC)	46
Chinese version(TC)	48

Introduction

GRAID SupremeRAID™ is the most powerful, high-speed data protection solution specially designed for NVMe SSDs. GRAID SupremeRAID™ installs a virtual NVMe controller onto the operating system and integrates a high-performance, AI processor equipped 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 Windows and how to manage the RAID components using the command-line interface.

GRAID SupremeRAID™ Specifications

GRAID SupremeRAID™ Driver Specifications	
Supported GRAID Models	SR-1000, SR-1010
Supported RAID levels	RAID 0, 1, 10, 5, 6
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	4
Maximum number of virtual drives per drive group	8
Maximum size of the drive group	Defined by the physical drive sizes

Supported Operating Systems

Operating System	x64
Windows	10
Windows Server	2019, 2022

Installation

This section describes installing the GRAID SupremeRAID™ software package for Windows.

Prerequisites

Before installing the software package, ensure that the system meets the following requirements:

1. Minimum system requirements:
 - CPU: 2 GHz or faster with at least 8 cores
 - RAM: 16 GB
 - An available PCIe Gen3 or Gen4 x16 slot
2. The GRAID SupremeRAID™ card is installed into a PCIe x16 slot
3. The IOMMU function is disabled in the system BIOS.
4. The GRAID SupremeRAID™ software package is downloaded from the GRAID, or GRAID partner, website.

Note: The GRAID SupremeRAID™ SR1000 and SR1010 drivers are different packages. To recognize your driver is supporting SR1000 or SR1010, please check the version code which is in the driver filename < graid-sr-x.x.x-xxx-xx.xxxxxxxx.**0x0**.x64 >, **000** is for SR1000, **010** is for SR1010.

Installing on Windows

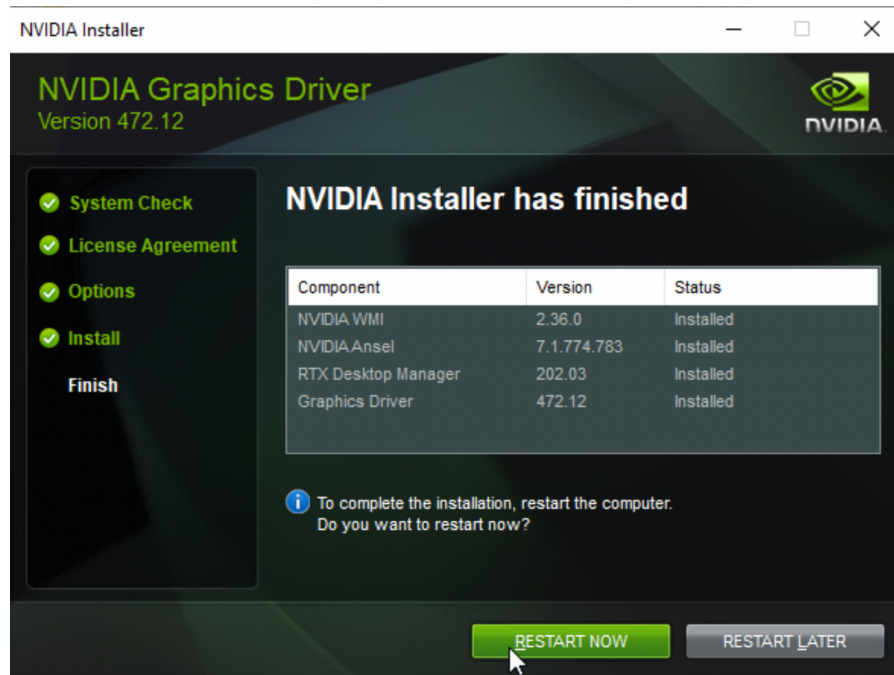
Note: You must install the NVIDIA driver and the Visual C++ Redistributable before installing the GRAID SupremeRAID™ driver for Windows.

To install the GRAID SupremeRAID™ driver on Windows 10, Windows Server 2019, or Windows Server 2022 systems:

1. Download the latest version of the NVIDIA driver, Visual C++ Redistributable, and the GRAID SupremeRAID™ driver.

Dependency	
NVIDIA Driver for Windows 10	511.79
Visual C++ Redistributable	VC_redist.x64.exe
Driver Package	
SR-1000 and SR1010	Download link

1. Install the NVIDIA driver and follow the instructions.



2. From the powershell or t running as Administrator, check the GPU Number and set the TCC mode for the GPU.

```
nvidia-smi
nvidia-smi -dm 1 -i [GPU_NUMBER]
```

Output example:

```
Windows PowerShell
Copyright (c) 2018 Microsoft Corporation. All rights reserved.

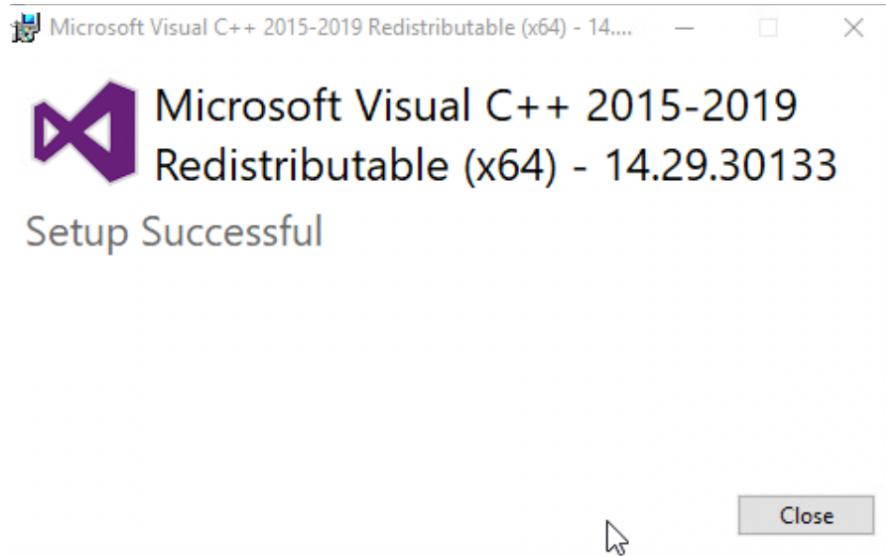
PS C:\Users\Administrator>nvidia-smi
Thu Nov 18 07:04:41 2021

+-----+
| NVIDIA-SMI 472.12      Driver Version: 472.12      CUDA Version: 11.4      |
+-----+-----+
| GPU   Name           TCC/WDDM  Bus-Id  Disp.A   Volatile Uncorr. ECC  |
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage   GPU-Util  Compute M.  |
|                                           MIG M.           |
+-----+-----+
|  0   NVIDIA T1000    TCC      00000000:0B:00_0 Off      539MiB / 3984MiB      0%      E. Process  |
| 40%   53C   P8     N/A / 50W      |                                     N/A      |
+-----+-----+

+-----+
| Processes: |
| GPU   GI   CI        PID   Type   Process name                      GPU Memory |
| ID   ID   ID             |              Usage                     |
+-----+-----+
|  0   N/A  N/A       3020   C     Insufficient Permissions           509MiB   |
+-----+

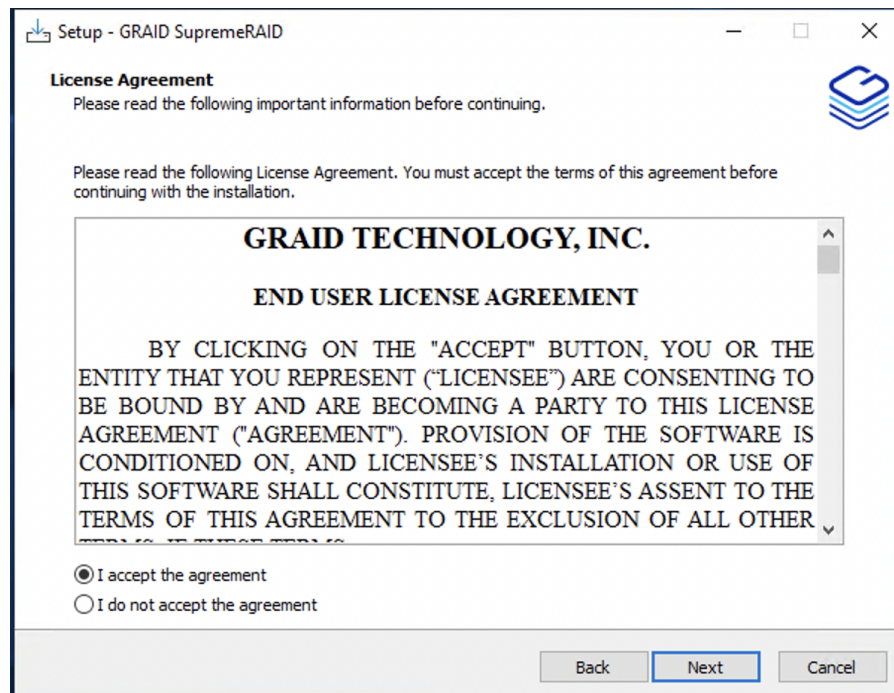
PS C:\Users\Administrator>nvidia-smi -dm 1 -i 0
Driver model is already set to TCC for GPU 00000000:0B:00.0.
All done.
```

4. Install the Visual C++ Redistributable and follow the instructions.

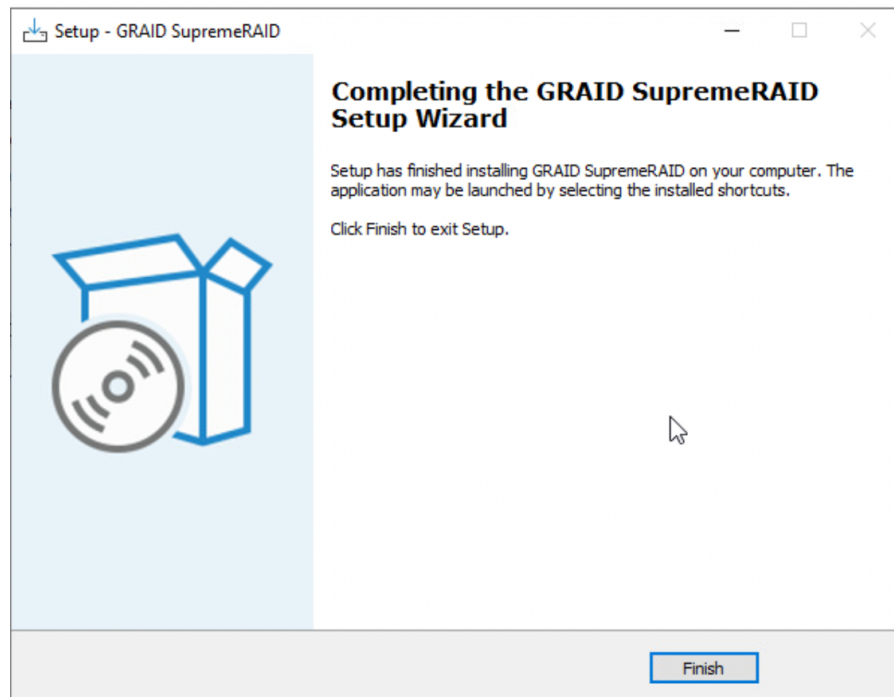


5. Install the GRAID SupremeRAID™ driver and follow the instructions.

1. In end-user license agreement page, you can scroll down the license content, after you review the license, select accept agreement or not and click "Next" or "Cancel" to proceed



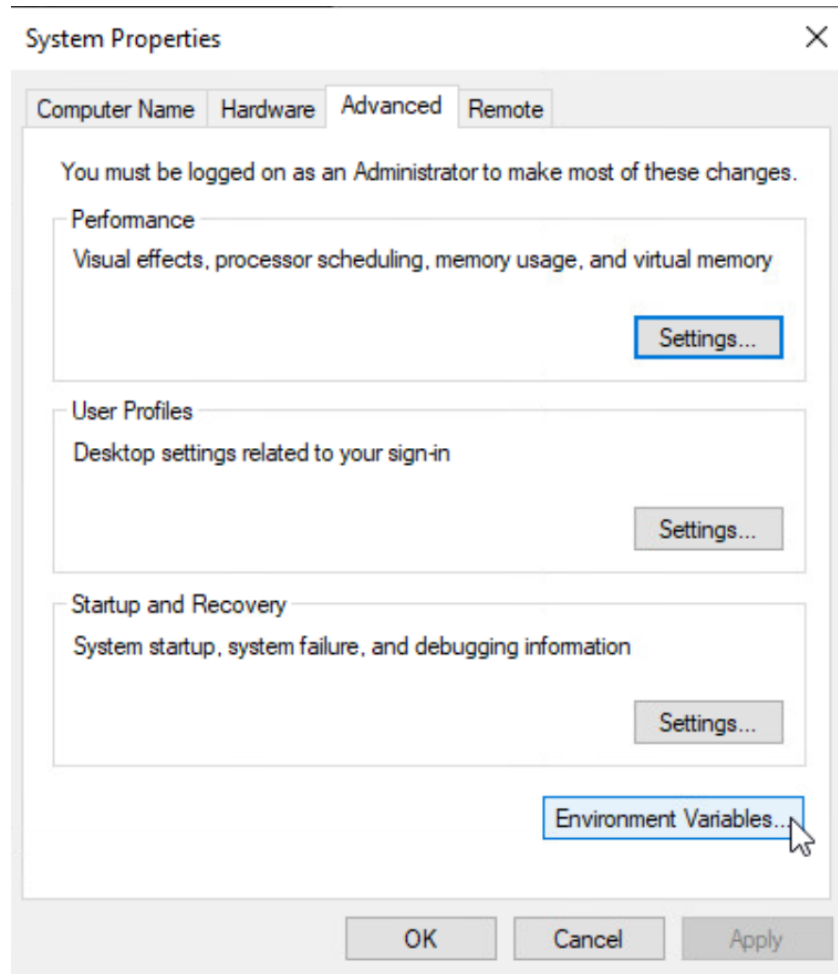
2. Install GRAID SupremeRAID driver page.



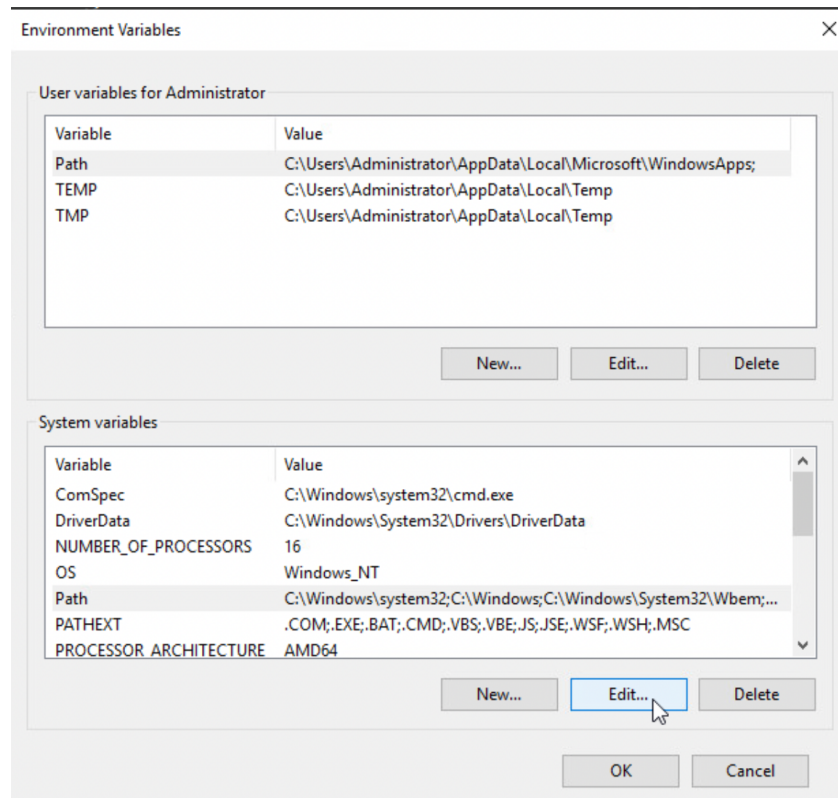
3. Execute the GRAID SupremeRAID™ driver. There are two ways to launch the GRAID service.

Appending the GRAID SupremeRAID™ driver folder path into the "Environment Variables"


1. Press "Windows Key + R" and type "sysdm.cpl" to open the **System Properties** dialog.
2. Select the **Advanced** tab and click **Environment Variables....**



3. From the **Environment Variables** dialog, select the **Path** row and click **Edit**.



4. Add the GRAID SupremeRAID™ driver path, save it, and reboot the system.




Changing directories to the GRAID SupremeRAID™ driver folder

1. From the powershell or command prompt running as administrator, change to the directory containing the GRAID SupremeRAID™ driver folder.

```
> cd "%PROGRAMFILES%\GRAID SupremeRAID\bin"
```

Output example:



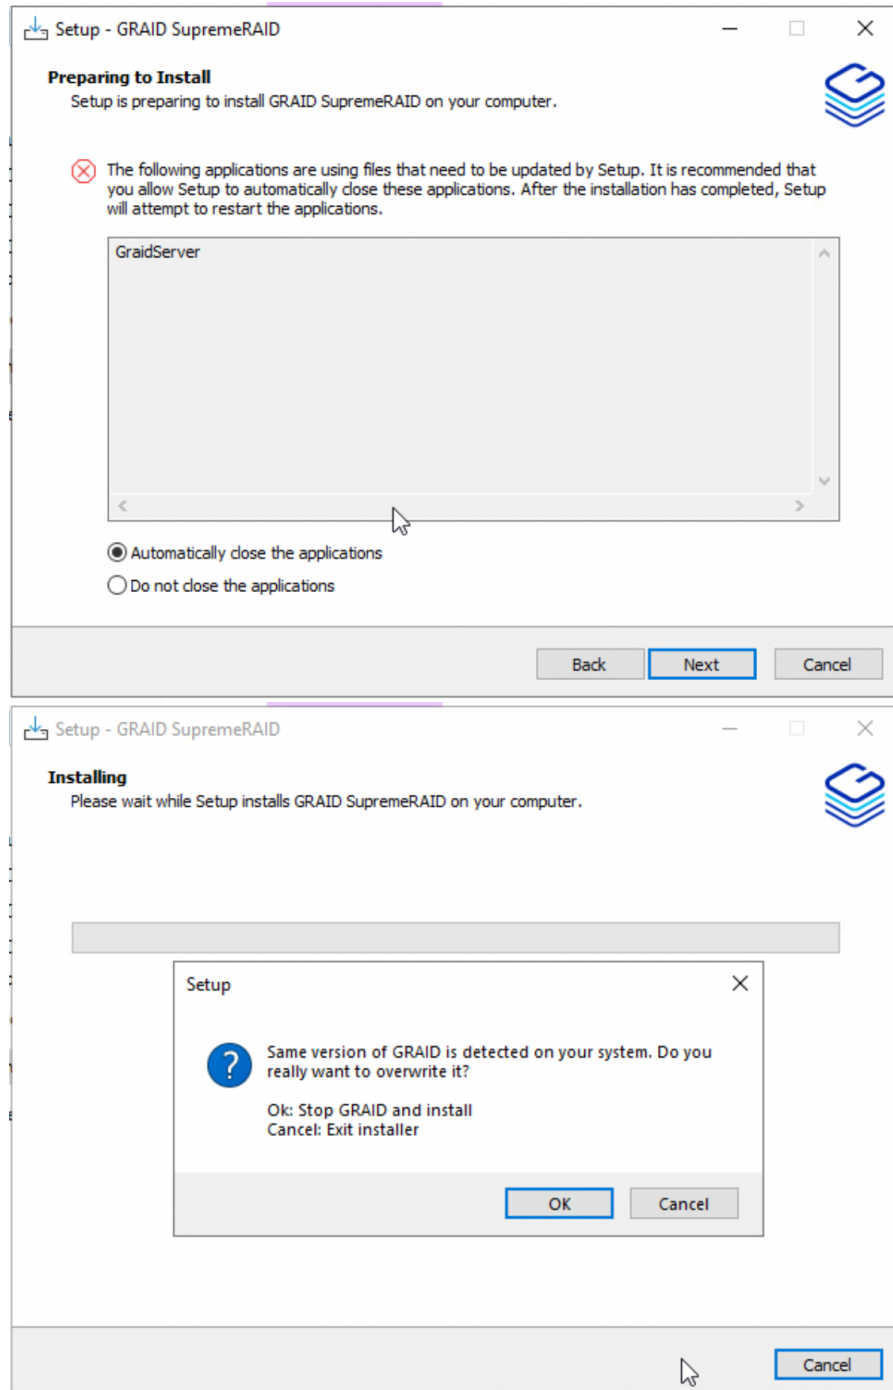
```
Windows PowerShell
Copyright (c) 2018 Microsoft Corporation. All rights reserved.

PS C:\Users\Administrator>cd "C:\Program Files\GRAID SupremeRAID\bin"
PS C:\Program Files\GRAID SupremeRAID\bin>
```

Upgrading the Software

To upgrade the software:

1. Download the newest GRAID SupremeRAID™ driver and execute the installer.
2. When the warning message appears, follow the instructions and uninstall the previous driver.



Management

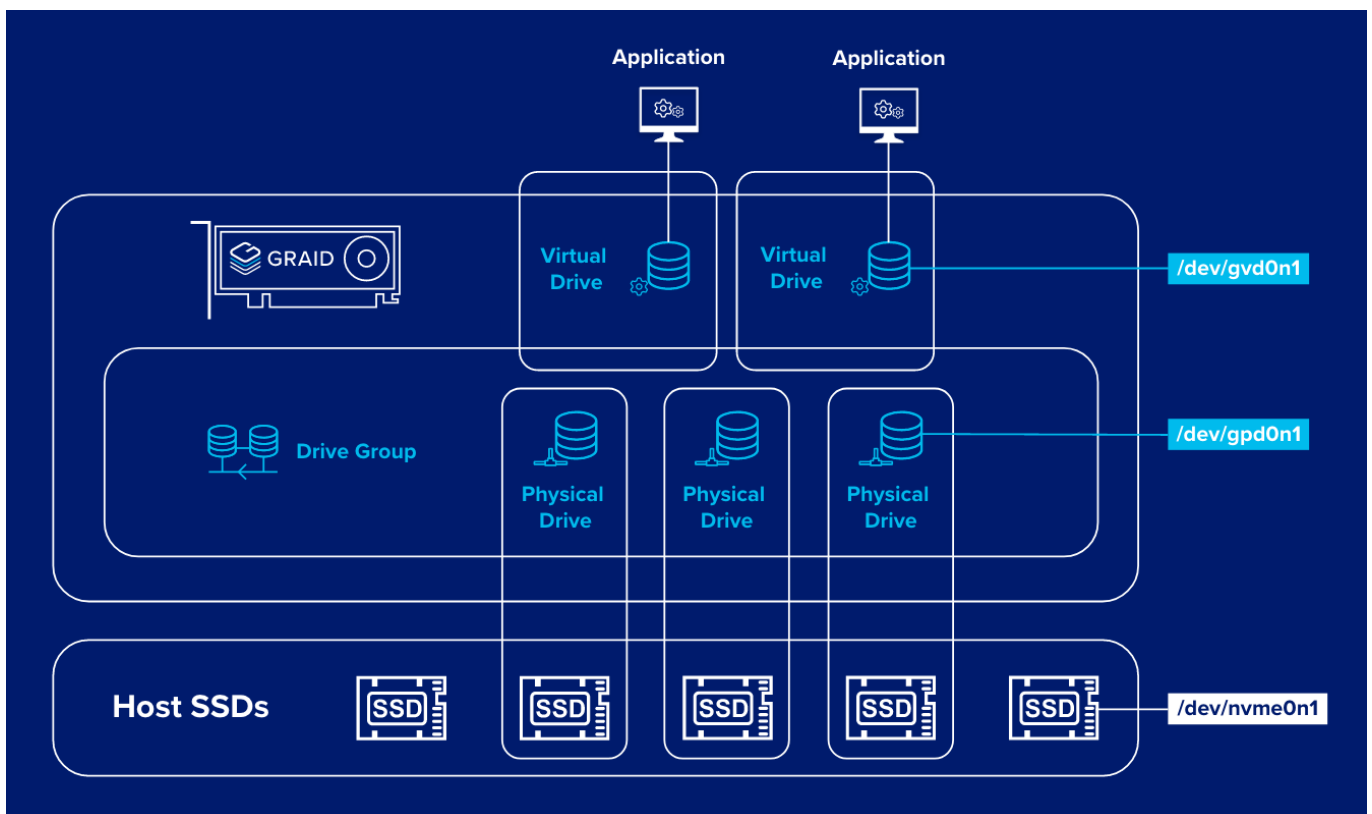
Overview of the GRAID SupremeRAID™ Software Module

There are two major components of GRAID SupremeRAID Software Module:

1. **graidctl** - The command-line management tool.
2. **graid_server** - The management daemon that handles requests from **graidctl** to control the driver.

RAID Components

There are three major RAID logical components in SupremeRAID™, the **Physical Drive (PD)**, the **Drive Group (DG)**, and the **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. Once an SSD has been created as a physical drive, the SupremeRAID™ driver unbonds the SSD from the operating system, meaning the device node (.\PHYSICALDRIVE x) will disappear and is no longer accessible. At the same time, a corresponding device node is created by the SupremeRAID™ driver. You can check the SSD information, such as SSD model or SMART logs, using this device node. To control and access the SSD using `nvmeXn1`, you must first delete the corresponding physical drive.

Currently, SupremeRAID™ supports a total of **32** NVMe SSD physical drives.

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 the parity are synchronized. There are two types of the initialization processes.

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

Currently, SupremeRAID™ supports a total of **4** 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 (`\\.\PHYSICALDRIVEx`) 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 **8** virtual drives in each drive group.

Overview of graidctl

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 you can use a dash to describe an OBJECT_ID range.

For example, to delete physical drives 1, 3, 4, and 5 simultaneously:

```
> graidctl delete physical_drive 1 3-5
```

- **flags**: Specifies optional flags.

For example:

- **--force**

Force to delete a physical drive

```
> graidctl delete physical_drive 0 --force
```

- **--json** (Supported in coming version)

Print output in json format. This option might assist with API implementation.

```
> graidctl list virtual_drive --format json
```

To get help, run *graidctl help* from the terminal window.

Managing Licenses

You can apply the license and check license information.

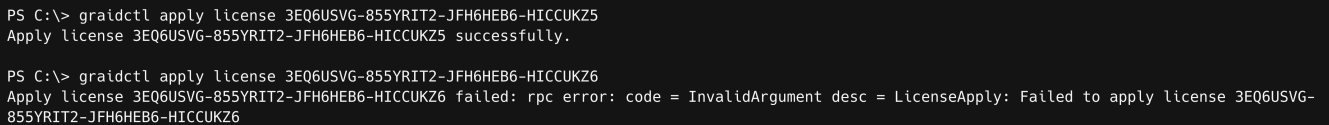
To complete the installation, apply the license.

Applying the License

To apply the license and complete the installation, run:

```
> graidctl apply license <LICENSE_KEY>
```

Output example for applying invalid license and valid license:



```
PS C:\> graidctl apply license 3EQ6USVG-855YRIT2-JFH6HEB6-HICCUKZ5
Apply license 3EQ6USVG-855YRIT2-JFH6HEB6-HICCUKZ5 successfully.

PS C:\> graidctl apply license 3EQ6USVG-855YRIT2-JFH6HEB6-HICCUKZ6
Apply license 3EQ6USVG-855YRIT2-JFH6HEB6-HICCUKZ6 failed: rpc error: code = InvalidArgument desc = LicenseApply: Failed to apply license 3EQ6USVG-855YRIT2-JFH6HEB6-HICCUKZ6
```

Note:

When applying the license, you might need to provide the serial number of the NVIDIA GPU to GRAID Technical Support.

To obtain NVIDIA GPU serial number, run:

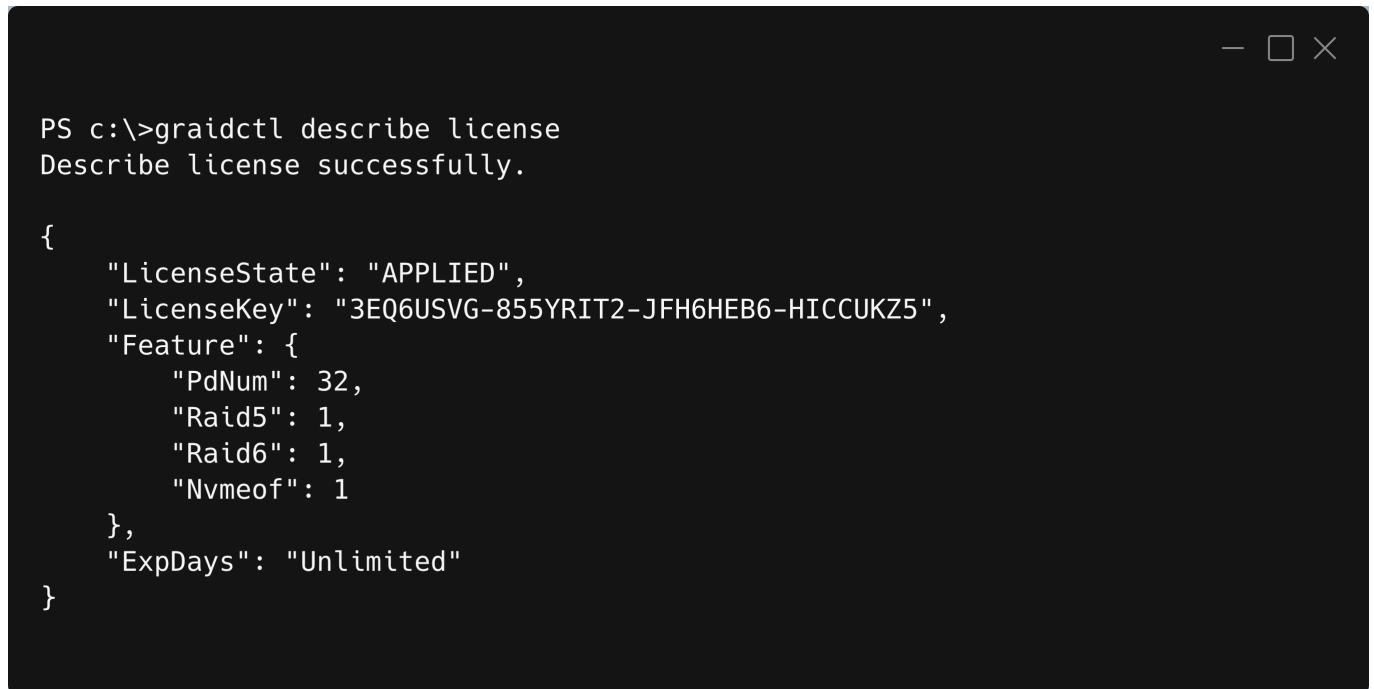
```
> nvidia-smi -q | findstr -i serial
```


Checking License Information

To obtain the license information, run:

```
> graidctl describe license
```

Output example:



```
PS c:\>graidctl describe license
Describe license successfully.

{
  "LicenseState": "APPLIED",
  "LicenseKey": "3EQ6USVG-855YRIT2-JFH6HEB6-HICCUKZ5",
  "Feature": {
    "PdNum": 32,
    "Raid5": 1,
    "Raid6": 1,
    "Nvmeof": 1
  },
  "ExpDays": "Unlimited"
}
```

Output content:

Field	Description
License State	The license state.
License Key	The applied license key.
Feature	The feature set of the license key.
ExpDays	The expiration date of the license key.

The license state:

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

Viewing Host Drive Information

Listing the NVMe Drives

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

```
> graidctl list nvme_drive
```

Output example:



```
PS C:\> graidctl list nvme_drive
List nvme drive successfully.
```

DEVICE PATH (4)	NQN	MODEL	CAPACITY
\\?\pci#ven_15ad&dev_07f0&subsys_07f015ad&rev_...	nqn.2014.08.org.nvmexpress:15ad15adVMWare NVME_0001	VMware Virtual NVMe Disk	16 GB
\\?\pci#ven_15ad&dev_07f0&subsys_07f015ad&rev_...	nqn.2014.08.org.nvmexpress:15ad15adVMWare NVME_0003	VMware Virtual NVMe Disk	16 GB
\\?\pci#ven_15ad&dev_07f0&subsys_07f015ad&rev_...	nqn.2014.08.org.nvmexpress:15ad15adVMWare NVME_0002	VMware Virtual NVMe Disk	16 GB
\\?\pci#ven_15ad&dev_07f0&subsys_07f015ad&rev_...	nqn.2014.08.org.nvmexpress:15ad15adVMWare NVME_0000	VMware Virtual NVMe Disk	16 GB

Output content:

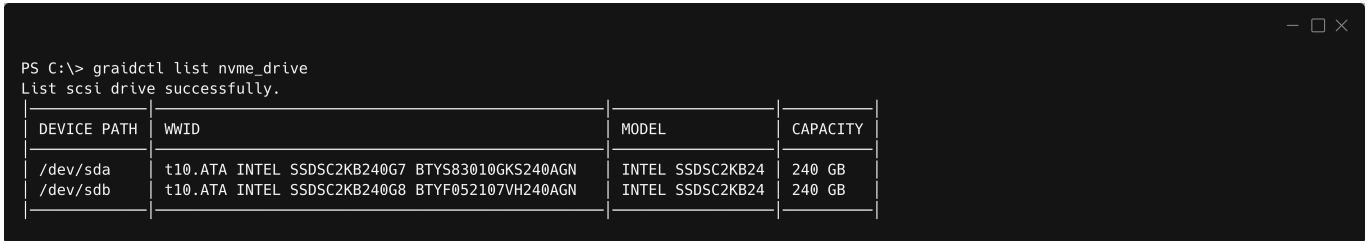
Field	Description
DEVICE PATH	This field displays the block device path of the drive.
NQN	This field displays the NVMe Qualified Name of the drive.
MODEL	This field displays the model number of the drive.
CAPACITY	This field displays the capacity of the drive.

Listing SAS/SATA Drives

To list all SAS/SATA drives that can be used as physical drives, run:

```
> graidctl list scsi_drive
```

Output example:



```
PS C:\> graidctl list nvme_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

Note: Due to a Windows limitation, SAS/SATA drives cannot currently be used to create physical drives.

Output content:

Field	Description
DEVICE PATH	This field displays the block device path for the drive.
WWID	This field displays the Worldwide Identification of the drive.
MODEL	This field displays the model number for the drive.
CAPACITY	This field displays the capacity of the drive.

Managing Physical Drives

Creating a Physical Drive

To create a physical drive, run:

```
> graidctl icreate physical_drive
```

Tip: The "icreate" command enables you select the specific drive to use to create the physical drive.

Note: Only NVMe drives can be used to create physical drives.

Output example for simultaneously creating multiple physical drives with the device path and NQN:

```
PS C:\> graidctl icreate physical_drive
? Choose one or more NVMe devices: [Use arrows to move, space to select, <right> to all, <left> to none, type to filter]
[x] Device Path:\\?\pci#ven_15ad&ddev_07f0&subsyst_07f015ad&rev_00#4&126f3afb&0&00b9#{2accfe60-c130-11d2-b082-00a0c91efb8b},
NQN/WWID:nqn.2014.08.org.nvmexpress:15ad15adVM
Ware NVME_0001, Model:VMware Virtual NVMe Disk , Capacity:16 GB
> [ ] Device Path:\\?\pci#ven_15ad&ddev_07f0&subsyst_07f015ad&rev_00#4&2f010251&0&0aa#{2accfe60-c130-11d2-b082-00a0c91efb8b},
NQN/WWID:nqn.2014.08.org.nvmexpress:15ad15adVMWare NVME_0003, Model:VMware Virtual NVMe Disk , Capacity:16 GB
[ ] Device Path:\\?\pci#ven_15ad&ddev_07f0&subsyst_07f015ad&rev_00#4&31783dac&0&00c1#{2accfe60-c130-11d2-b082-00a0c91efb8b},
NQN/WWID:nqn.2014.08.org.nvmexpress:15ad15adVM
Ware NVME_0002, Model:VMware Virtual NVMe Disk , Capacity:16 GB
[ ] Device Path:\\?\pci#ven_15ad&ddev_07f0&subsyst_07f015ad&rev_00#4&c28b61f&0&00b1#{2accfe60-c130-11d2-b082-00a0c91efb8b},
NQN/WWID:nqn.2014.08.org.nvmexpress:15ad15adVMW
are NVME_0000, Model:VMware Virtual NVMe Disk , Capacity:16 GB
-----
? Choose one or more NVMe devices: Device Path:\\?\pci#ven_15ad&ddev_07f0&subsyst_07f015ad&rev_00#4&126f3afb&0&00b9#{2accfe60-c130-11d2-b082-
00a0c91efb8b}, NQN/WWID:nqn.2014.
08.org.nvmexpress:15ad15adVMWare NVME_0001, Model:VMware Virtual NVMe Disk , Capacity:16 GB, Device Path:\\?
\pci#ven_15ad&ddev_07f0&subsyst_07f015ad&rev_00#4&2
f010251&0&0aa#{2accfe60-c130-11d2-b082-00a0c91efb8b}, NQN/WWID:nqn.2014.08.org.nvmexpress:15ad15adVMWare NVME_0003, Model:VMware Virtual NVMe Disk
, Capacit
y:16 GB, Device Path:\\?\pci#ven_15ad&ddev_07f0&subsyst_07f015ad&rev_00#4&31783dac&0&00c1#{2accfe60-c130-11d2-b082-00a0c91efb8b},
NQN/WWID:nqn.2014.08.org.nvmexpress:15ad15ad
VMWare NVME_0002, Model:VMware Virtual NVMe Disk , Capacity:16 GB, Device Path:\\?
\pci#ven_15ad&ddev_07f0&subsyst_07f015ad&rev_00#4&c28b61f&0&00b1#{2accfe60-c1
30-11d2-b082-00a0c91efb8b}, NQN/WWID:nqn.2014.08.org.nvmexpress:15ad15adVMWare NVME_0000, Model:VMware Virtual NVMe Disk
,
Capacity:16 GB
? Create new physical drive now? (y/N) y
-----
? Create new physical drive now? Yes
Icreate physical drive PD0 (nqn.2014.08.org.nvmexpress:15ad15adVMWare NVME_0001) successfully.

Icreate physical drive PD1 (nqn.2014.08.org.nvmexpress:15ad15adVMWare NVME_0003) successfully.

Icreate physical drive PD2 (nqn.2014.08.org.nvmexpress:15ad15adVMWare NVME_0002) successfully.

Icreate physical drive PD3 (nqn.2014.08.org.nvmexpress:15ad15adVMWare NVME_0000) successfully.
```

Note:

Ensure that the system or other applications are not on the physical drive before creating or replacing it.

Listing the Physical Drives

To list all of the physical drives, run:

```
> graidctl list physical_drive
```

Output example:

```
PS C:\> graidctl list physical_drive
List physical drive successfully.
```

PD ID (4)	DG ID	NQN/WWID	MODEL	CAPACITY	SLOT ID	STATE
0	N/A	nqn.2014.08.org.nvmexpress:15ad15adVMWare NVME_0001	VMware Virtual NVMe Disk	16 GB	N/A	UNCONFIGURED_GOOD
1	N/A	nqn.2014.08.org.nvmexpress:15ad15adVMWare NVME_0003	VMware Virtual NVMe Disk	16 GB	N/A	UNCONFIGURED_GOOD
2	N/A	nqn.2014.08.org.nvmexpress:15ad15adVMWare NVME_0002	VMware Virtual NVMe Disk	16 GB	N/A	UNCONFIGURED_GOOD
3	N/A	nqn.2014.08.org.nvmexpress:15ad15adVMWare NVME_0000	VMware Virtual NVMe Disk	16 GB	N/A	UNCONFIGURED_GOOD

Output content:

Field	Description
SLOT ID	This field displays the slot ID of the corresponding NVMe/SAS/SATA drive. Note that the PD ID is not related to the SLOT ID, and that you must set the physical drives using the PD ID.
DG ID	This field displays the drive group ID of the physical drive.
PD ID	This field displays the 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 ID is used for any subsequent operations.
NQN/WWID	This field displays the NQN or WWID of corresponding NVMe/SAS/SATA drive.
MODEL	This field displays the model number of the corresponding NVMe/SAS/SATA drive.
CAPACITY	This field displays the capacity of corresponding NVMe/SAS/SATA drive.
STATE	This field displays the physical drive state.

Physical drive STATE:

STATE	Description
ONLINE	The physical drive was added to a drive group and is ready to work.
HOTSPARE	The physical drive is configured as hot spare drive.
FAILED	The physical drive is detected but it is not functioning normally.
OFFLINE	The physical drive is marked as offline.
REBUILD	The physical drive is being rebuilt.

STATE	Description
MISSING	The physical drive cannot be detected.
INCONSISTENT	The data in the physical drive is inconsistent. This condition generally occurs when the physical drive is in the REBUILD state and the system encounters an abnormal crash.
UNCONFIGURED_GOOD	The physical drive did not join a drive group.
UNCONFIGURED_BAD	The physical drive did not join any drive group and is not functioning normally.

Deleting Physical Drives

To delete physical drives, run:

```
> graidctl delete physical_drive <PD_ID>
```

Output example for deleting multiple physical drives simultaneously:

```
PS C:\> graidctl delete physical_drive 2
Delete physical drive PD2 failed: rpc error: code = NotFound desc = PD2 is still using by DG0
PS C:\> graidctl delete physical_drive 0 1 5-7
Delete physical drive PD0 successfully.
Delete physical drive PD1 successfully.
Delete physical drive PD5 successfully.
Delete physical drive PD6 successfully.
Delete physical drive PD7 successfully.
```

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

Describing a Physical Drive

To check the detailed information for a physical drive, run:

```
> graidctl describe physical_drive <PD_ID>
```

Output example:

```
PS C:\> graidctl describe physical_drive 0
Describe physical drive PD0 successfully.
{
  "SlotId": 0,
  "DgId": -1,
  "PdId": 0,
  "Guid": "nqn.2019-10.com.kioxia:KCM61VUL3T20:Z080A038T1L8",
  "Model": "KCM61VUL3T20",
  "Capacity": "3.2 TB",
  "State": "UNCONFIGURED_GOOD",
  "Attrs": {
    "hotspare": "",
    "locating": "false"
  }
}
```

Locating a Physical Drive

To locate a physical drive, run:

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

To stop locating a physical drive, run:

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

Marking a Physical Drive Online or Offline

To mark a physical drive as online or offline, run:

```
> 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, run:

```
> graidctl edit physical_drive <PD_ID> hotspare global
```

To assign a physical drive as the hot spare for a specific drive group, run:

```
> 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.

Replacing a Nearly Worn-Out or Broken SSD

1. Mark the physical drive as bad using the following command. (You can skip this step if the physical drive is in the **MISSING** or other abnormal state.)

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

2. Replace the NVMe SSD. The state of the prior physical drive will indicate **FAILED**.
3. Check the NQN of the new SSD.

```
> graidctl list nvme_drive
```

4. Create a new physical drive on the new SSD.

```
> graidctl create physical_drive <NEW_SSD_NQN>
```

5. Replace the physical drive.

```
> graidctl replace physical_drive <OLD_PD_ID> <NEW_PD_ID>
```

6. Delete the old physical drive

```
> graidctl delete physical_drive <OLD_PD_ID>
```


Managing Drive Groups

Creating a Drive Group

To create a drive group, run:

```
> graidctl create drive_group <RAID_MODE> (PD_IDs) [--background-init]
```

Output example showing three drive groups were created:

```
PS C:\> graidctl create drive_group raid1 0-1
Create drive group DG0 successfully.
PS C:\> graidctl create drive_group raid5 2-4
Create drive group DG1 successfully.
PS C:\> graidctl create drive_group raid6 5-9
Create drive group DG2 successfully.
```

Required parameters:

Option	Description
RAID_MODE	This field specifies the RAID mode of the drive group. Entries must be all uppercase or all lowercase. (For example, RAID6 or raid6 are both correct)
PD_IDs	This field specifies the ID of the physical drive joining the drive group.

Optional parameters:

Option	Description
	Default option.
--	Use standard methods to initialize the drive group.
background-init	When all the physical drives in the drive group support the deallocate dataset management command, it is used to synchronize the data, or parity, between the physical drives during the creation of the drive group.

Important Note:

Wait for the drive group initialization to complete. DO NOT power-off or reboot the system when the drive_group state is INIT/RESYNC/RECOVERY.

Use the command below to check drive_group state:

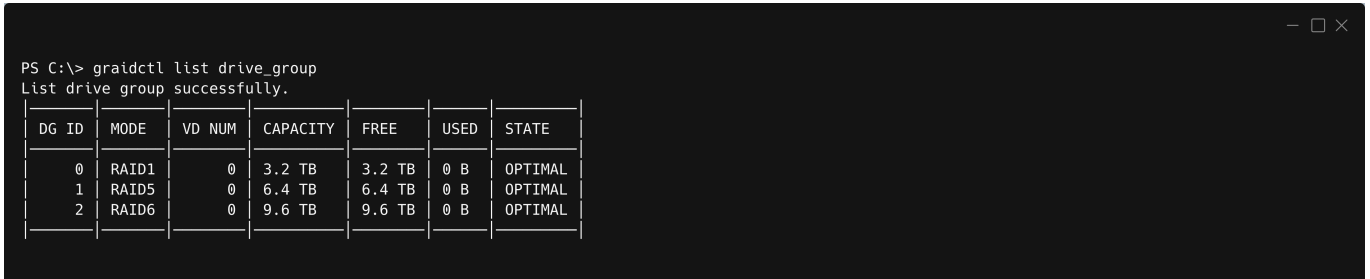
```
> graidctl list drive_group
```

Listing Drive Groups

To list all drive groups, run:

```
> graidctl list drive_group
```

Output example:



```
PS C:\> graidctl list drive_group
List drive group successfully.
```

DG ID	MODE	VD NUM	CAPACITY	FREE	USED	STATE
0	RAID1	0	3.2 TB	3.2 TB	0 B	OPTIMAL
1	RAID5	0	6.4 TB	6.4 TB	0 B	OPTIMAL
2	RAID6	0	9.6 TB	9.6 TB	0 B	OPTIMAL

Output content:

Field	Description
DG ID	This field displays the drive group ID.
MODE	This field displays the drive group RAID mode.
VD NUM	This field displays the number of virtual drives in the drive group.
CAPACITY	This field displays the total usable capacity of the drive group.
FREE	This field displays the unused space of the drive group.
USED	This field displays the used space of the drive group.
STATE	This field displays the drive group state.

Drive Group STATE:

STATE	Description
OFFLINE	The drive group does not function normally. This condition is usually caused when the number of damaged physical drives exceeds the limit.
OPTIMAL	The drive group is in optimal state.
DEGRADED	The drive group is available and ready, but the number of missing or failed physical drives has reached the limit.
PARTIALLY_DEGRADED	The drive group is available and ready for use, but some physical drives are missing or failed.
RECOVERY	The drive group is recovering.

STATE	Description
FAILED	The drive group does not function normally.
INIT	The drive group is initializing.
RESYNC	The drive group is re-synchronizing. 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.
RESCUE	The drive group is in rescue mode.

Deleting Drive Groups

To delete drive groups, run:

```
> graidctl delete drive_group <DG_ID>
```

```
PS C:\> graidctl delete drive_group 1
Delete drive group DG1 failed: rpc error: code = FailedPrecondition desc = DG1 still has IVD(s)
PS C:\> graidctl delete drive_group 0 2
Delete drive group DG0 successfully.
Delete drive group DG2 successfully.
```

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

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

Describing a Drive Group

To display detailed information for a drive group, run:

```
> graidctl describe drive_group <DG_ID>
```

Output example:

```
PS C:\> graidctl describe drive_group 1
Describe drive group DG1 successfully.
{
  "DgId": 1,
  "Mode": "RAID5",
  "Capacity": "6.4 TB",
  "Free": "0 B",
  "Used": "6.4 TB",
  "State": "OPTIMAL",
  "PdIds": [
    2,
    3,
    4
  ],
  "VdNum": 1,
  "Attrs": {
    "rebuild_speed": "low"
  }
}
```

To set the rebuild speed for a drive group, run:

```
> 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 the drive group, run:

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

To stop locating all the physical drives in drive group, run:

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

Degradation and Recovery

- When multiple drive groups require simultaneous recovery, the drive groups recover individually.
- When multiple physical drives in the same drive group require rebuilding, the physical drives are rebuilt simultaneously.

Rescue Mode

When a damaged drive group is initialized, or when 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 the rescue mode, run:

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

Managing Virtual Drives

Creating a Virtual Drive

To create a virtual drive, run:

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

Output example:

```
PS C:\> graidctl create virtual_drive 0
Create virtual drive VD0 in DG0 successfully.
PS C:\> graidctl create virtual_drive 1 100G
Create virtual drive VD0 in DG1 successfully.
PS C:\> graidctl create virtual_drive 2 1T
Create virtual drive VD0 in DG2 successfully.
```

Listing Virtual Drives

To list virtual drives, run:

```
> graidctl list virtual_drive [--dg-id=<DG_ID>] [--vd-id=<VD_ID>]
```

Output example:

```
PS C:\> graidctl list virtual_drive
List virtual drive successfully.
```

VD ID (4)	DG ID	SIZE	DEVICE PATH	STATE
0	0	3.2 TB	\\.\PHYSICAKDRIVE1	OPTIMAL
0	1	100 GB	\\.\PHYSICAKDRIVE2	OPTIMAL
0	2	1.0 TB	\\.\PHYSICAKDRIVE3	OPTIMAL
1	2	1.0 TB	\\.\PHYSICAKDRIVE4	OPTIMAL

Output content:

Field	Description
DG ID	This field displays the drive group ID.
VD ID	This field displays the virtual drive ID.
SIZE	This field displays the usable size of the virtual drive.
DEVICE PATH	This field displays the device path of the virtual drive.
NQN	This field displays the NQN of the virtual drive.
STATE	This field displays the virtual drive state. It is identical to the drive group state.

Field	Description
EXPORTED	This field displays 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/gvd0n) is created.

Virtual Drive STATE:

Identical to the drive group state.

STATE	Description
OFFLINE	The drive group does not function normally. This condition is usually caused when the number of damaged physical drives exceeds the limit.
OPTIMAL	The drive group is in optimal state.
DEGRADED	The drive group is available and ready, but the number of missing or failed physical drives has reached the limit.
PARTIALLY_DEGRADED	The drive group is available and ready for use, but some physical drives are missing or failed.
RECOVERY	The drive group is recovering.
FAILED	The drive group does not function normally.
INIT	The drive group is initializing.
RESYNC	The drive group is re-synchronizing. 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.
RESCUE	The drive group is in rescue mode.

Deleting Virtual Drives

To delete virtual drives, run:

```
> graidctl delete virtual_drive <DG_ID> <VD_ID> [--force]
```

Output example:

```
PS C:\> graidctl delete virtual_drive 0 0
Delete virtual drive VD0 from DG0 successfully.
PS C:\> graidctl delete virtual_drive 2 0-1
Delete virtual drive VD1 from DG2 successfully.
Delete virtual drive VD0 from DG2 successfully.
```

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

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

1. Create a physical drive

```
> graidctl icreate physical_drive
Create physical drive PD0 successfully.
Create physical drive PD1 successfully.
Create physical drive PD2 successfully.
Create physical drive PD3 successfully.
Create physical drive PD4 successfully.
```

2. Create a drive group.

```
> graidctl create drive_group raid5 0-4
Create drive group DG0 successfully.
```

3. Create a virtual drive.

```
> graidctl create virtual_drive 0
Create virtual drive VD0 successfully.
```

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

```
> graidctl list virtual_drive --dg_id=0
```

DG ID	VD ID	SIZE	DEVICE PATH	STATE
0	0	500 GB	\\.\PHYSICAKDRIVE1	OPTIMAL

Basic Troubleshooting

Sequential Read Performance is not as Expected on a New Drive Group

Unlike SAS/SATA hard drives, many NVMe SSDs support the deallocate 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.

But for other SSDs, the performance is not as expected when reading unwritten sectors after issuing the deallocate 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 GRAID 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 that SupremeRAID™ can control the SSD. During the unbinding process, when the APST feature is enabled, the NVMe driver attempts 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 functioning normally.

Installer Error Message: "... nvml.dll not found..." Appears When Installing the SupremeRAID™ Driver

Error message:

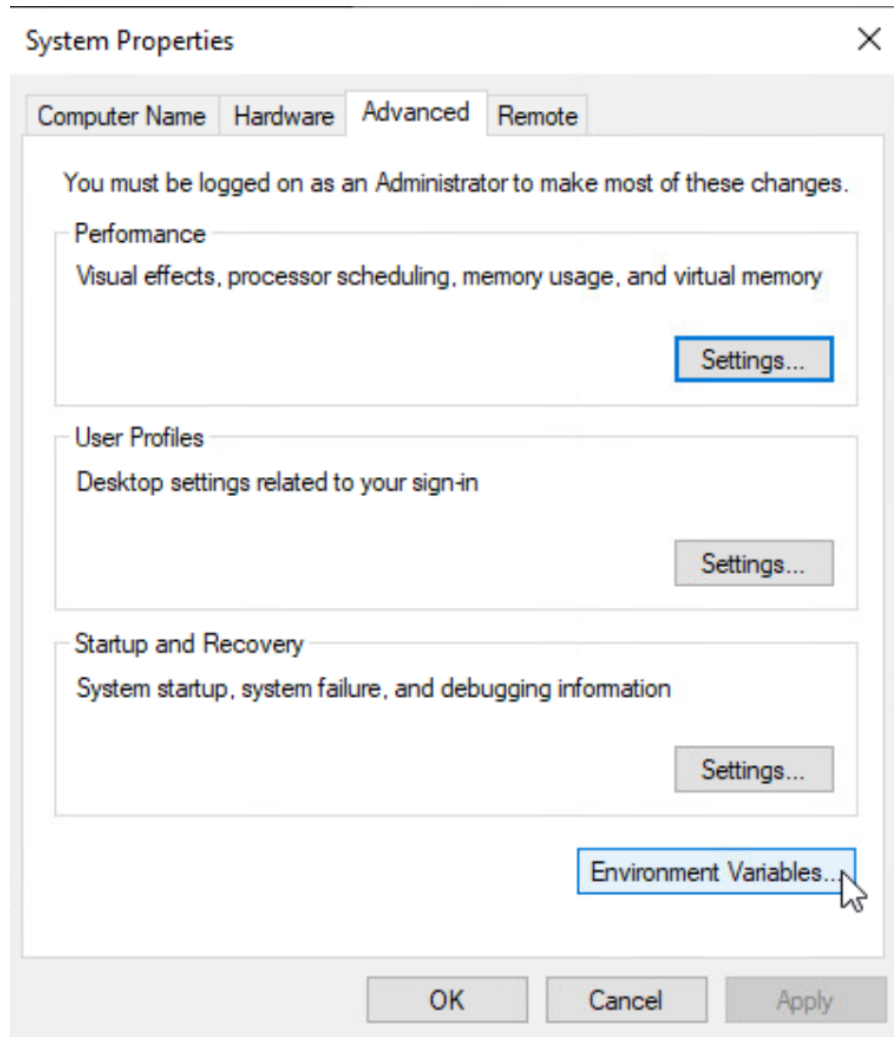


During installation the SupremeRAID™ driver must detect the presence of the NVIDIA driver. However, the NVIDIA driver does not automatically append the installation folder to the environment variables during its installation. This condition prevents the SupremeRAID™ driver from detecting the NVIDIA driver and causes the error message. Appending the NVSMI folder path to the environment variables fixes the issue.

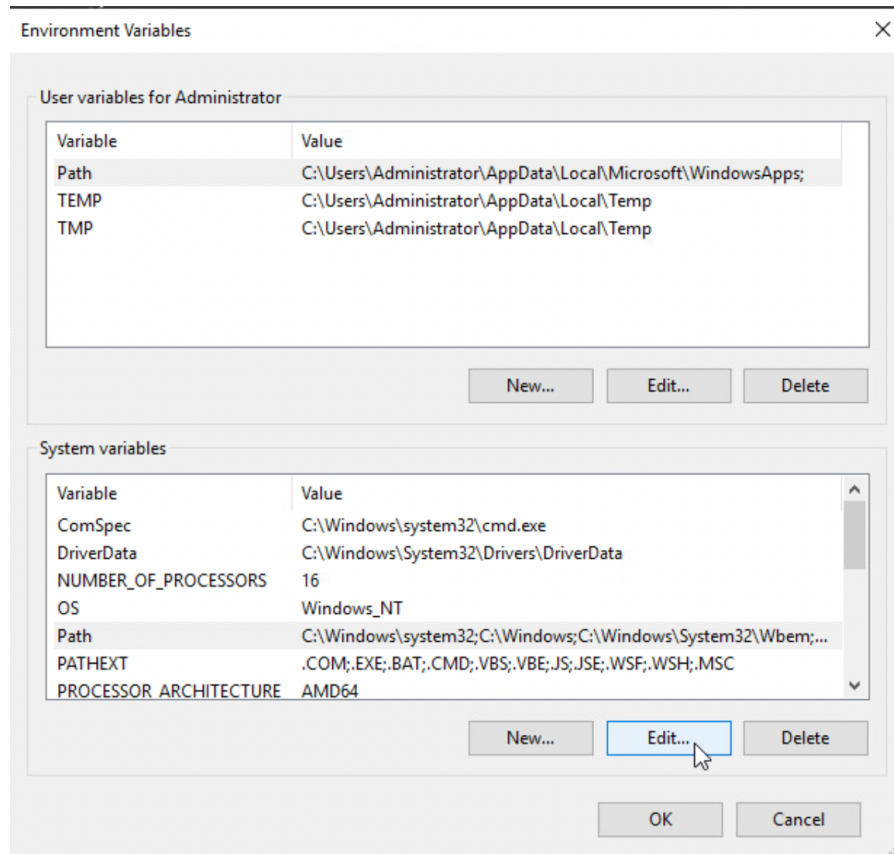
To append the NVSMI folder path to the environment variables:

1. Press "Windows Key + R" and type "sysdm.cpl" to open the **System Properties** dialog.

2. Select the **Advanced** tab and click **Environment Variables...**

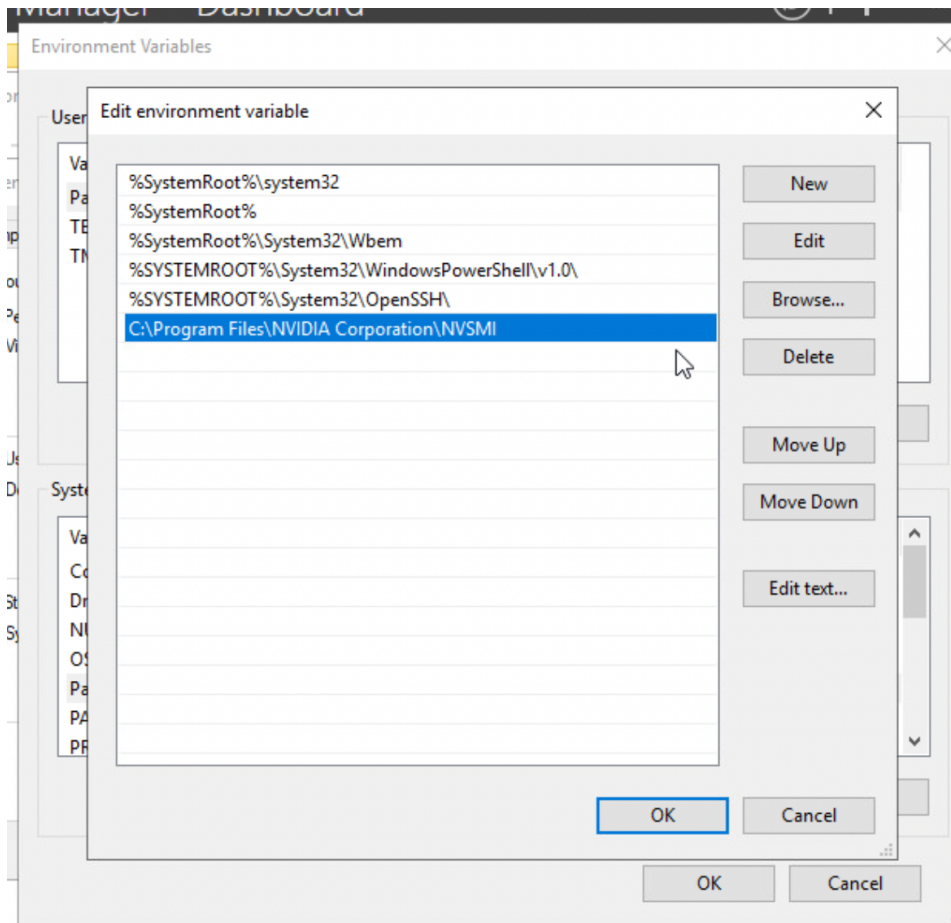


3. From the **Environment Variables** dialog, select the **Path** row and click **Edit**.



4. Add the NVSMI path, save it, and reboot the system.

NVSMI path:
 %PROGRAMFILES%\NVIDIA Corporation\NVSMI\



Error Message: "GRAID service is not running"

Use either of the following methods to check the status of the GRAID service.

Check the GRAID service using a command prompt or PowerShell

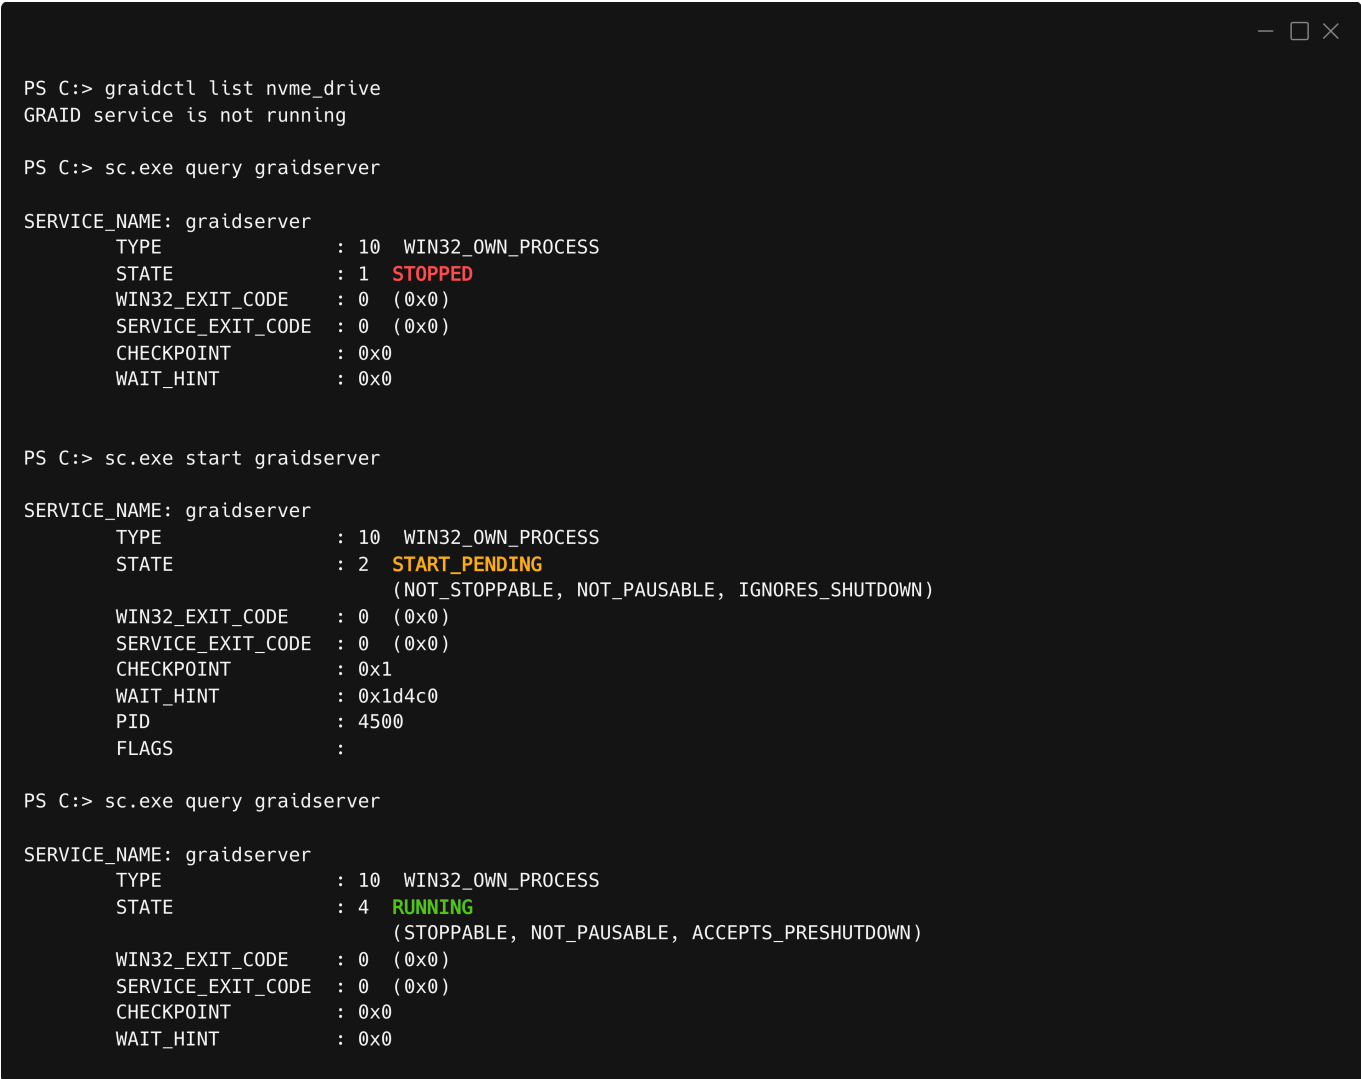
1. From a command prompt or PowerShell, run:

```
sc.exe query graidserver
```

2. If the service is not "RUNNING", type the following command and wait a moment before checking its status again.

```
sc.exe start graidserver
```

Output example:



```
PS C:> graidctl list nvme_drive
GRAID service is not running

PS C:> sc.exe query graidserver

SERVICE_NAME: graidserver
        TYPE               : 10  WIN32_OWN_PROCESS
        STATE                : 1   STOPPED
        WIN32_EXIT_CODE       : 0   (0x0)
        SERVICE_EXIT_CODE    : 0   (0x0)
        CHECKPOINT           : 0x0
        WAIT_HINT            : 0x0

PS C:> sc.exe start graidserver

SERVICE_NAME: graidserver
        TYPE               : 10  WIN32_OWN_PROCESS
        STATE                : 2   START_PENDING
                        (NOT_STOPPABLE, NOT_PAUSABLE, IGNORES_SHUTDOWN)
        WIN32_EXIT_CODE       : 0   (0x0)
        SERVICE_EXIT_CODE    : 0   (0x0)
        CHECKPOINT           : 0x1
        WAIT_HINT            : 0x1d4c0
        PID                 : 4500
        FLAGS                 :

PS C:> sc.exe query graidserver

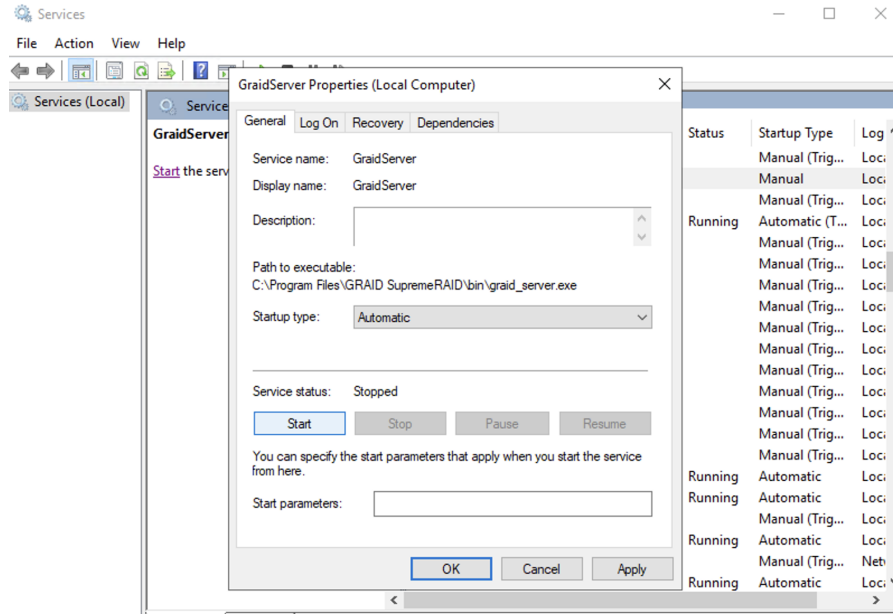
SERVICE_NAME: graidserver
        TYPE               : 10  WIN32_OWN_PROCESS
        STATE                : 4   RUNNING
                        (STOPPABLE, NOT_PAUSABLE, ACCEPTS_PRESHUTDOWN)
        WIN32_EXIT_CODE       : 0   (0x0)
        SERVICE_EXIT_CODE    : 0   (0x0)
        CHECKPOINT           : 0x0
        WAIT_HINT            : 0x0
```

Check the GRAID service using Windows Service Manager

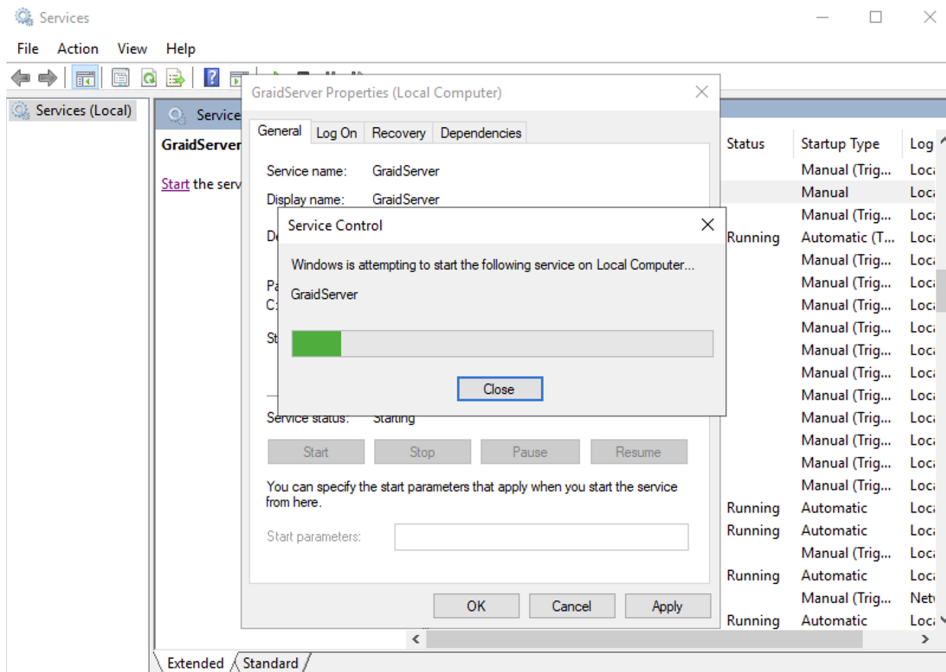
1. Open **Windows Service Manager** and search "graidserver" to check its status.
2. If the service is not "RUNNING", click **Start** and wait a moment for the process to complete.

Output example:

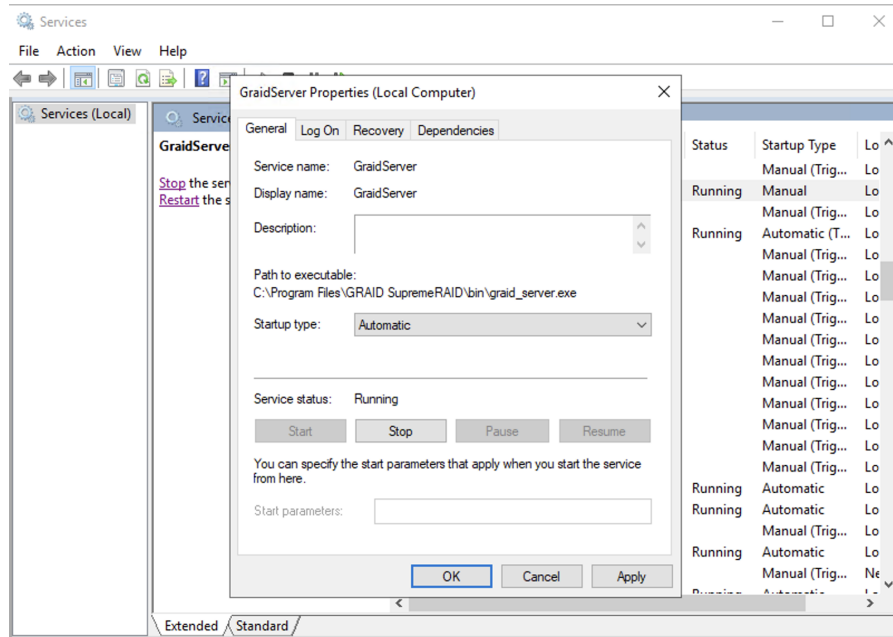
- `graid_server` is not running



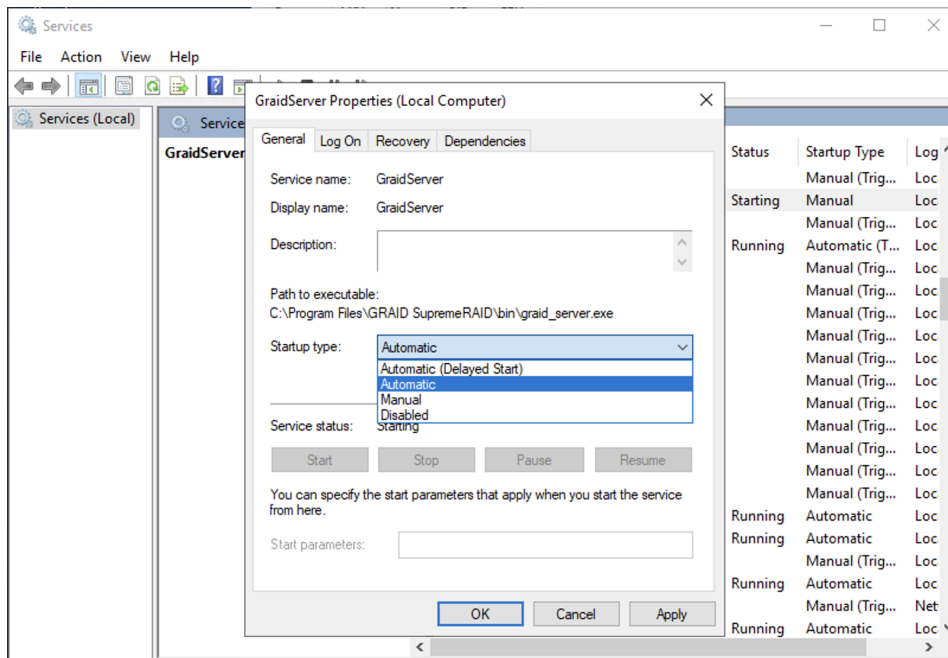
- starting `graid_server`



- `graid_server` is running



Note: From the **GraidServer Properties** dialog, ensure that the **Startup type:** option is set to **Automatic**.



Different LED Blink Patterns on the Backplane

You might notice that the HDD/SSD activity indicator blink pattern is different on GRAID SupremeRAID™ than on traditional RAID cards.

GRAID SupremeRAID™ does not require a buffering or caching mechanism to improve read/write performance like traditional RAID cards. This feature causes GRAID SupremeRAID™ indicators to blink differently than traditional RAID cards.

Appendix

Manually Migrating the RAID Configuration Between Hosts

To manually migrate the RAID configuration between hosts:

1. Periodically backup the configuration file **%PROGRAMFILES%\GRAID SupremeRAID\confgraid.conf** from the original host. To move the configuration file to another system.
2. Setup the target host and ensure that the GRAID service is stopped.

Note

When the target host already contains an installed and running SupremeRAID™ card, stop and restart the service using the **graid.conf** file from the original system.

3. Move all SSDs from the original host to the new host.
4. Copy the configuration backup file to the new host using the same path.
5. Start the GRAID service directly if the original card also moved to the new host.

```
> sc start graidservice
```

Otherwise, you have to apply the new license of the new GRAID SupremeRAID card on the new host

```
> graidctl apply license <LICENSE_KEY>
```

Enabling Virtual Machines with GPU Passthrough

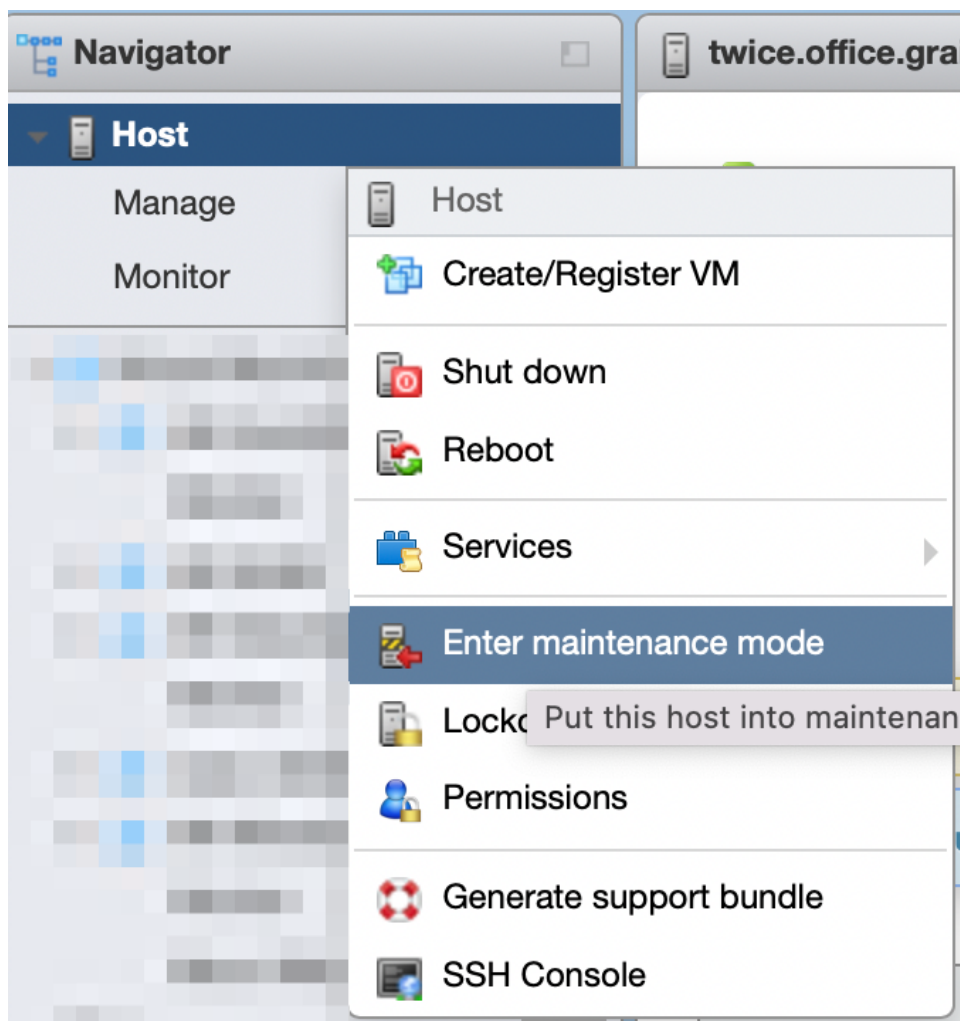
You can create virtual machines with GRAID SupremeRAID™ support to maximize performance.

Hypervisor support: VMWare ESXi 7.0U3

Configuring Hosts for NVIDIA GPU Device Passthrough

Put ESXi host into maintenance mode

1. From the **Navigator** menu, select **Host > Enter maintenance mode**.



Managing PCI Device Passthrough

1. From the **Navigator** menu, select **Manage > Hardware > PCI Devices**. The **Passthrough Configuration** page appears listing all of the available pass-through devices.
2. Select the NVIDIA T1000 (Quadro T1000 Mobile) and its audio device.
3. Click **Toggle passthrough**.
4. Check the Passthrough status. It should be **Active**.

System: Hardware Licensing Packages Services Security & users					
PCI Devices <input type="checkbox"/> Toggle passthrough <input type="checkbox"/> Configure SR-IOV <input type="checkbox"/> Hardware label <input type="checkbox"/> Reboot host <input type="checkbox"/> Refresh <input type="text" value="Search"/> 					
Power Management					
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		
<input type="checkbox"/> 0000:40:04.0	Advanced Micro Devices, Inc. [AMD] Starship/Matisse PCIe Dummy Host Bridge	Not capable	Not capable		
<input type="checkbox"/> 0000:40:05.0	Advanced Micro Devices, Inc. [AMD] Starship/Matisse PCIe Dummy Host Bridge	Not capable	Not capable		
<input type="checkbox"/> 0000:40:07.0	Advanced Micro Devices, Inc. [AMD] Starship/Matisse PCIe Dummy Host Bridge	Not capable	Not capable		
<input type="checkbox"/> 0000:40:07.1	Advanced Micro Devices, Inc. [AMD] Starship/Matisse Internal PCIe GPP Bridg...	Not capable	Not capable		
123 items					

Configuring Virtual Machines

Attach PCI devices to virtual machines

1. From the **Edit VM setting** page, choose **Virtual Hardware > Add other device > PCI device**.
2. Choose Quadro T1000 and its Audio device as the two PCI devices.

▶ PCI device 1	TU117GLM [Quadro T1000 Mobile] - 0000:42:00.0	<input type="checkbox"/>
▶ PCI device 2	<class> Audio device - 0000:42:00.1	<input type="checkbox"/>

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.

3. Choose **Virtual Hardware > Memory**.
4. Check **Reserve all guest memory (All locked)**.

Virtual Hardware

VM Options

Add hard disk
 Add network adapter
 Add other device

▶ CPU	8	▼		
▼ Memory				
RAM	16		GB	▼
Reservation	16384	▼	MB	▼
<input checked="" type="checkbox"/> Reserve all guest memory (All locked)				

Enabling Point-to-Point (P2P) on the Virtual Machine for Best Performance

- From the **Edit VM setting** page, choose **VM Options** > **Advanced** > **Configuration Parameters** > **Edit Configuration...**

Virtual Hardware

VM Options

▶ General Options	VM Name: <input type="text" value="tiff-Ubuntu"/>
▶ VMware Remote Console Options	<input type="checkbox"/> Lock the guest operating system when the last remote user disconnects
▶ VMware Tools	Expand for VMware Tools settings
▶ Power management	Expand for power management settings
▶ Boot Options	Expand for boot options
▼ Advanced	
Settings	<input type="checkbox"/> Disable acceleration <input checked="" type="checkbox"/> Enable logging
Debugging and statistics	Run normally ▼
Swap file location	<input checked="" type="radio"/> Default Use the settings of the cluster or host containing the virtual machine. <input type="radio"/> Virtual machine directory Store the swap file in the same directory as the virtual machine. <input type="radio"/> Datastore specified by host Store the swap files in the datastore specified by the host to be used for swap files. If not possible, store the swap files in the same directory as the virtual machine. Using a datastore that is not visible to both hosts during vMotion might affect the vMotion performance for the affected virtual machines.
Configuration Parameters	<input type="button" value="Edit Configuration..."/>

Add two parameters:

```
hypervisor.cpuid.v0 = "FALSE"
pciPassthru.allowP2P = "TRUE"
```

Compatible NVMe Drives List

The following NVMe drives passed GRAID qualification and can be used with GRAID SupremeRAID™.

GRAID updates this list when new NVMe drives pass the qualification process.

Manufacturer	Series	Interface	FormFactor
Dapustor	H3200	PCIe Gen3 x 4	2.5 inch U.2
Intel	DC P4510	PCIe Gen3 x 4	2.5 inch U.2
Intel	DC P4610	PCIe Gen3 x 4	2.5 inch U.2
Intel	D7-P5316	PCIe Gen4 x 4	2.5 inch U.2
Intel	D7-P5510	PCIe Gen4 x 4	2.5 inch U.2
Intel	Optane™ P5800X	PCIe Gen4 x 4	2.5 inch U.2
Kioxia	CD5	PCIe Gen3 x 4	2.5 inch U.2
Kioxia	CD6	PCIe Gen4 x 4	2.5 inch U.3
Kioxia	CM6	PCIe Gen4 x 4	2.5 inch U.3
Memblaze	P6536	PCIe Gen4 x 4	2.5 inch U.2
Micron	9300	PCIe Gen3 x 4	2.5 inch U.2
Micron	7300	PCIe Gen4 x 4	2.5 inch U.2
Micron	7400	PCIe Gen4 x 4	2.5 inch U.3
Netlist	N1951	PCIe Gen3 x 4	2.5 inch U.2
Samsung	PM963	PCIe Gen3 x 4	2.5 inch U.2
Samsung	PM983	PCIe Gen3 x 4	2.5 inch U.2
Samsung	PM9A3	PCIe Gen4 x 4	2.5 inch U.2
Samsung	PM1733	PCIe Gen4 x 4	2.5 inch U.2
Western Digital Technologies	SN640	PCIe Gen3 x 4	2.5 inch U.2

For the latest information, see the [Compatible NVMe Drives List](#) on the GRAID website.

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 70950-1/62368-1 based safety standards around the world e.q. 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 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 burn. Allow to cool before servicing.


Chinese version(SC)

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！表面高温。接触可能导致灼伤。请再冷却后再使用。



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）

部件名称	有害物质						
	铅	汞	镉	六价铬	多溴联苯	多溴二苯醚	
PCB板	○	○	○	○	○	○	○
被动电子零件	X	○	○	○	○	○	○
主动电子零件	X	○	○	○	○	○	○
处理器	○	○	○	○	○	○	○
内存	○	○	○	○	○	○	○
结构间以及风扇	X	○	○	○	○	○	○
线材/连接器	X	○	○	○	○	○	○
焊接金属	○	○	○	○	○	○	○
助焊剂，锡膏，标签及耗材	○	○	○	○	○	○	○

本表格依据SJ/T 11364-2014的规定编制

○：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。

X：表示该有害物质至少在该部位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。

此表中所有名称中含“X” 的部件均符合RoHS立法。

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


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	○	○	○	○	○	○
結構間以及風扇	-	○	○	○	○	○
連結器	-	○	○	○	○	○
被動電子零件	-	○	○	○	○	○
主動電子零件	-	○	○	○	○	○
內存	○	○	○	○	○	○
線材	○	○	○	○	○	○
焊接金屬	○	○	○	○	○	○
助焊劑、錫膏、標籤及耗材	○	○	○	○	○	○

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外項目。

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