

SupremeRAID™ SE User Guide for Linux

Mar 2025



Copyright © 2021–2025 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: Mar 1, 2025

Table of Contents

- INTRODUCTION 5
 - SupremeRAID™ SE Specifications 6
 - RAID Components 7
 - Physical Drive (PD) 8
 - Drive Group (DG) 8
 - Virtual Drive (VD) 8
 - Controller (CX) 9
 - Limitations of SupremeRAID™ SE 9
- SIGN UP FOR THE SUPREMER RAID™ SE 9
- INSTALLATION 9
 - Prerequisites 11
 - BIOS setting 11
 - Installing the Software Driver 12
 - Using the Pre-installer and installer 12
- MANAGING PHYSICAL DRIVES 16
 - Creating a Physical Drive 17
 - Deleting a Physical Drive 18
 - Marking a Physical Drive Online/Offline/Good/Bad 18
 - Assigning a Hot Spare Drive 19
- MANAGING DRIVE GROUPS 20
 - Creating a Drive Group 20
 - Deleting a Drive Group 21
 - Degradation and Recovery 22
 - Rescue Mode 22

- MANAGING VIRTUAL DRIVES..... 23
 - Creating a Virtual Drive.....23
 - Deleting a Virtual Drive24

- FEATURES OVERVIEW 25
 - Consistency Checks25
 - Drive Copyback25
 - Double Failure Protection with Distributed Journaling.....26
 - Dataset Management (DSM) Deallocate command on virtual drives27

- ATTACHMENTS..... 28
 - Events for SupremeRAID™ SE28

INTRODUCTION

SupremeRAID™ SE is the most powerful, high-speed data protection solution specially designed for NVMe SSDs. SupremeRAID™ SE 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™ SE software package for Linux and how to manage the RAID components using the Graphical Management Console.

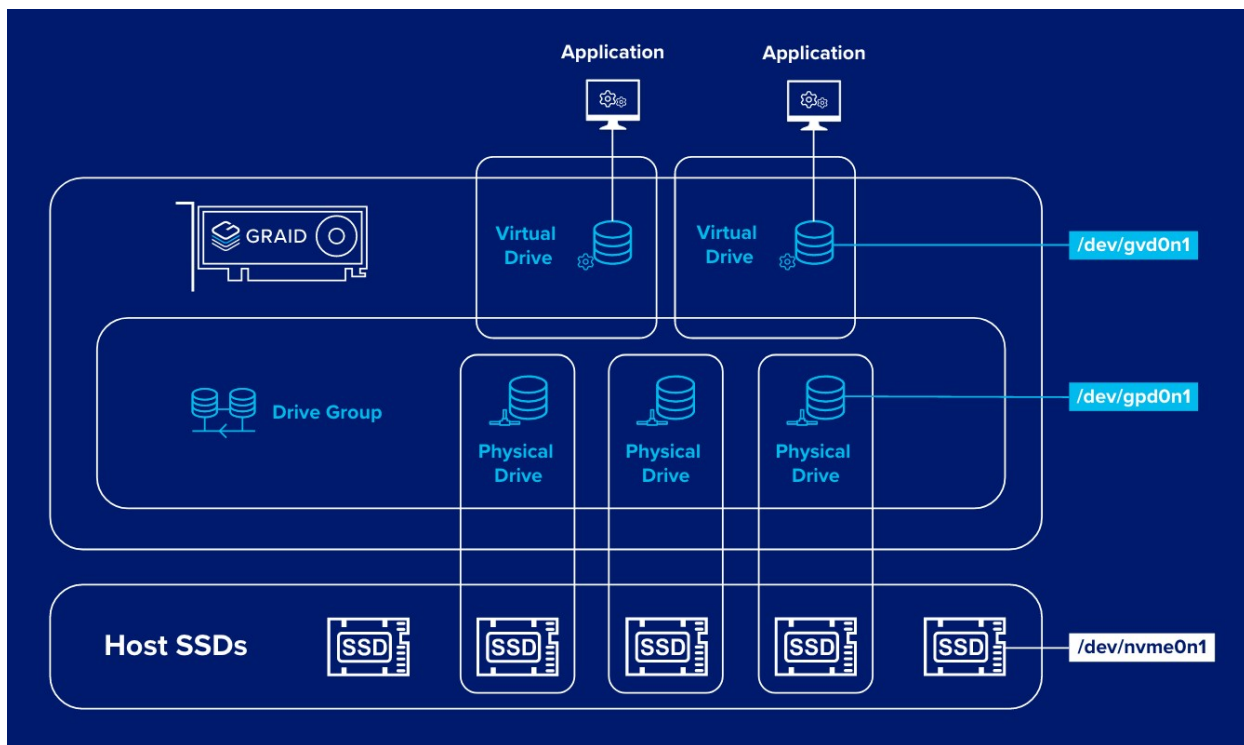
SUPREMER RAID™ SE SPECIFICATIONS

SupremeRAID™ SE Driver Specifications	
Supported Models	SE-TU, SE-AM, SE-AD, SE-BW
Supported RAID levels	RAID 0, 1, 5, 10,
Supported OS	<ul style="list-style-type: none"> ● AlmaLinux 8.5-8.10 (Kernel 4.18), 9.0-9.5(Kernel 5.14) ● CentOS 7.9 (Kernel 3.10 and 4.18), 8.3, 8.4, 8.5 (Kernel 4.18) ● Debian 11.6 (Kernel 5.10), 12 (Kernel 6.1) ● openSUSE Leap 15.2-15.3 (Kernel 5.3), 15.4-15.5 (Kernel 5.14) ● Oracle Linux 8.7-8.10 (RHCK 4.18 and UEK 5.15), 9.1-9.4 (RHCK 5.14 and UEK 5.15) ● Proxmox VE 8.1 (Kernel 6.5), 8.2 (Kernel 6.8) ● RHEL 7.9 (Kernel 3.10 or 4.18), 8.3-8.10 (Kernel 4.18), 9.0-9.5 (Kernel 5.14) ● Rocky Linux 8.5-8.10(Kernel 4.18), 9.0-9.5(Kernel 5.14) ● SLES 15 SP2-SP3(Kernel 5.3), 15 SP4-SP5 (Kernel 5.14) ● Ubuntu 20.04.0-20.04.4 (Kernel 5.4 and 5.15), 22.04.0-22.04.2 (Kernel 5.15), 22.04.3-22.04.4 (Kernel 6.2), 24.04 (Kernel 6.8)
Supported GPU models	Turing, Ampere, Ada Lovelace, Blackwell series
Recommended minimum drive number for each RAID level	RAID 0: at least one drives RAID 1: at least two drives RAID 5: at least three drives RAID 10: at least two drives
Maximum number of physical drives	8
Maximum number of drive groups	8
Maximum number of virtual drives per drive group	1023
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

There are four major RAID logical components in SupremeRAID™ SE:

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



Physical Drive (PD)

Since NVMe drives are not directly attached to the SupremeRAID™ SE controller, you must tell the controller which SSDs can be managed. After an SSD is created as a physical drive, the SupremeRAID™ SE 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™ SE 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™ SE supports a maximum 8 physical drives in one Drive Group.

Drive Group (DG)

The main component of RAID logic is a RAID group. When the drive group is created, the SupremeRAID™ SE 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 the physical drives in the drive group (DG) support the deallocate dataset management command, the SupremeRAID™ SE 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.

SupremeRAID™ SE supports a maximum of 8 Drive Groups.

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 run application directly on this device node

SupremeRAID™ SE supports a maximum of 1023 Virtual Drives in each Drive Group.

Controller (CX)


The controller is the core component of the RAID system. It provides detailed hardware information such as GPU serial number, temperature, and fan speed. RAID management relies on the controller, so the controller's state directly affects the underlying drive group operations.

Limitations of SupremeRAID™ SE

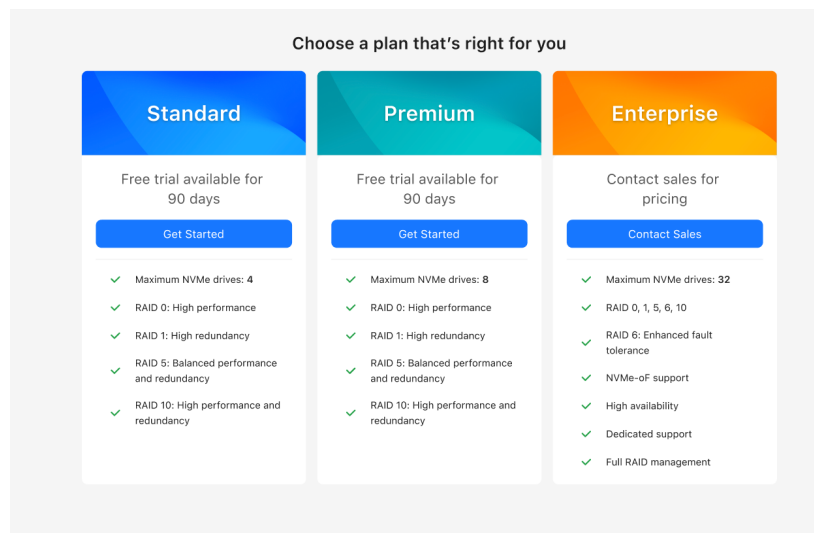
- System suspension and hibernation are currently unsupported due to a limitation in the NVIDIA driver.
- NVMe-oF drives are not supported in the Linux driver.
- If you have previously installed other SupremeRAID™ products, executing the SupremeRAID™ SE installation will prevent the restoration of previous product settings.
- After your SupremeRAID™ SE subscription expires, you will still be able to access and read your data, but the data cannot be written.

SIGN UP FOR THE SUPREME RAID™ SE

Step 1 You can sign up for the SupremeRAID™ SE through our [official website](#). Please create an account on the SupremeRAID™ SE Portal.



Step 2 Choose a subscription plan that is right for you.



Step 3 Upon successful subscription, please follow the steps in the next section to proceed with the installation.

INSTALLATION

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

Prerequisites

Before installing the software package, ensure that 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 the [Supported Operating Systems](#) section on our website.
 - An available PCIe Gen3 or Gen4 x16 slot
- The GPU card must be installed into a PCIe x16 slot.
- The SupremeRAID™ SE 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 SupremeRAID™ SE driver. The Installer contains the SupremeRAID™ SE driver package and will automatically detect your Linux distributions and install the appropriate files.
- Make sure a SupremeRAID™ SE compatible SSD drive is being used. SupremeRAID™ SE supports all form factors of NVMe drives, including M.2, U.2, U.3, EDSFF, and AIC.

BIOS Setting

- [OPTIONAL] The IOMMU function (AMD) or VT-d function (Intel) is recommended to be disabled in the system BIOS, typically found on the BIOS Advanced page.
- [OPTIONAL] 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.

Installing the Software Driver

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

Using the Pre-installer and Installer

The SupremeRAID™ SE 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™ SE driver on every supported Linux distribution. Use the following steps to prepare the environment and install the SupremeRAID™ SE 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. Please download the package in [Dependencies and Utilities](#).

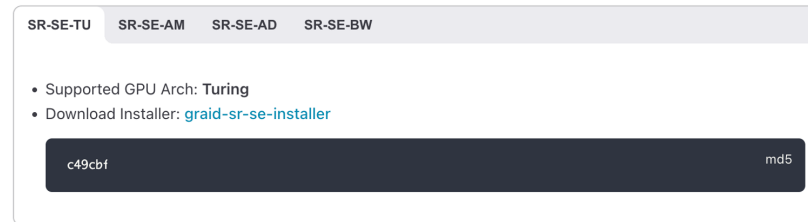
Dependencies and Utilities

	Links
NVIDIA Driver	NVIDIA-Linux-
SupremeRAID™ SE Pre-installer	graid-sr-se-pre-installer

```
$ sudo wget [filelink]
$ sudo chmod +x [Filename]
$ sudo ./[filename]
```

Step 2 Download the latest version of the installer and make it executable. Please download the installer from the [Driver Package](#) section.

Driver Package



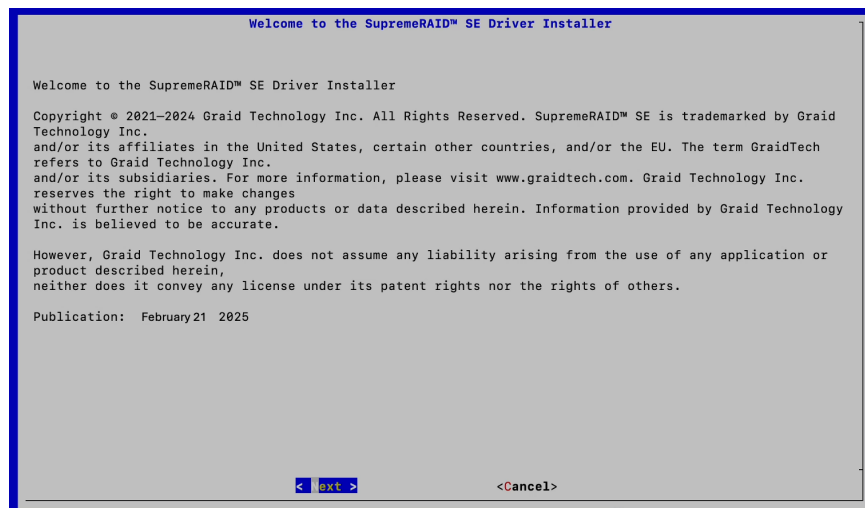
SR-SE-TU SR-SE-AM SR-SE-AD SR-SE-BW

- Supported GPU Arch: Turing
- Download Installer: [graid-sr-se-installer](#)

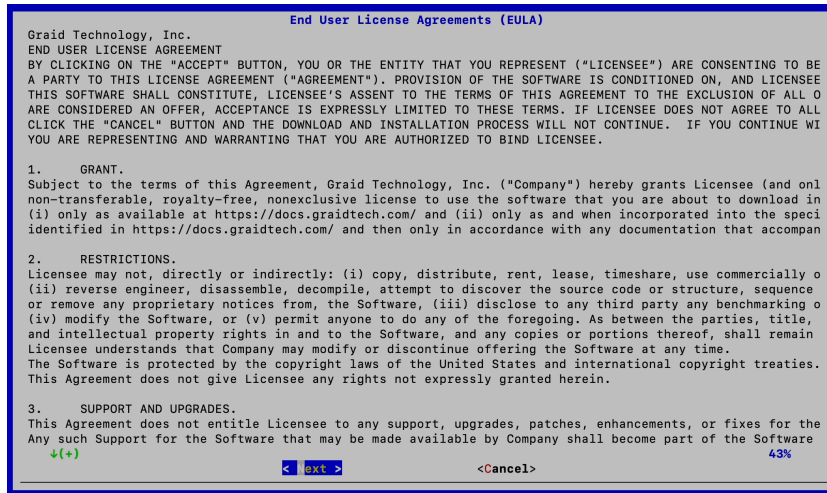
c49cbf md5

```
$ sudo wget[filelink]
$ sudo chmod +x [Filename]
$ sudo ./[filename]
```

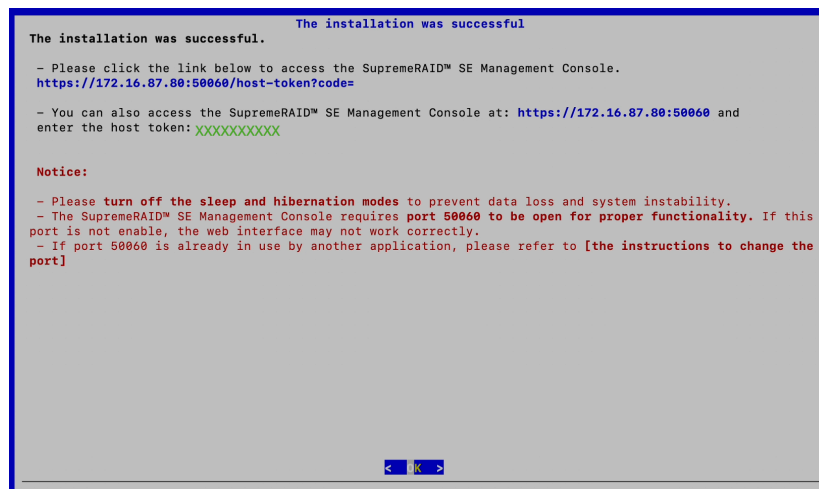
Step 3 At the Welcome page, select Next and press Enter to view the end-user license agreement.



Step 4 In the end-user license agreement page, you can scroll down to read the license content. After reviewing the license, accept the agreement, and click Next to proceed.



Step 5 Complete the installation. You can access the SupremeRAID™ SE Management Console. The SupremeRAID™ SE Management Console requires port 50060 to be open for proper functionality. If this port is not enabled, the web interface may not work correctly.

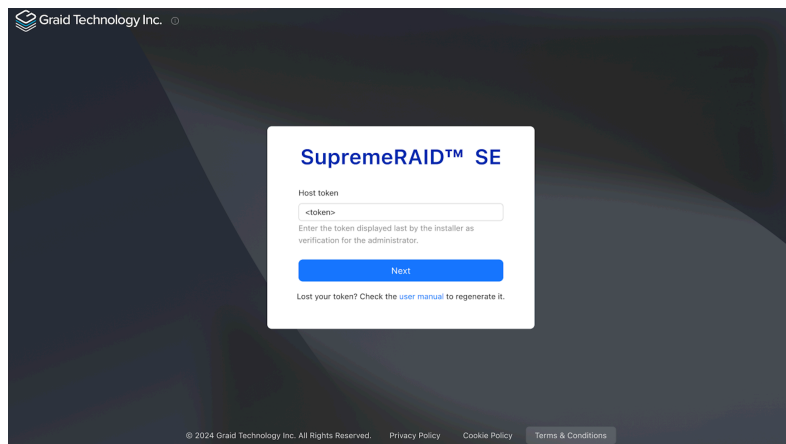


If port 50060 is occupied by another application, you can set up your own port and IP, please edit the configuration file `/etc/graidmgr/service.conf`.

For example, if you want to set the port and IP to 8888 and 123.456.7.889 respectively, it would look as follows:

```
[common]
web_port=8888
web_addr=123.456.7.889
```


Step 6 Enter the token and log in to the SupremeRAID™ SE Management Console.



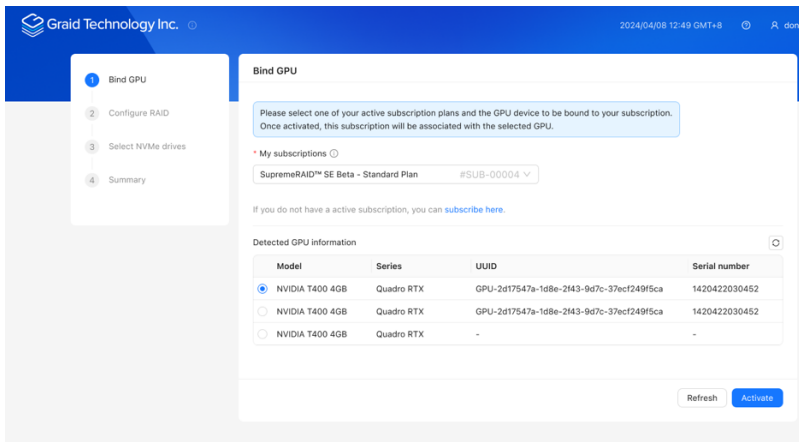
If you have lost the token, please use the following command to retrieve the host token.

```
sudo systemctl stop graid-mgr
sudo graid-mgr host_token gen
sudo systemctl start graid-mgr
```

Log in to the SupremeRAID™ SE Management Console using the email and password you [signed up](#) with.



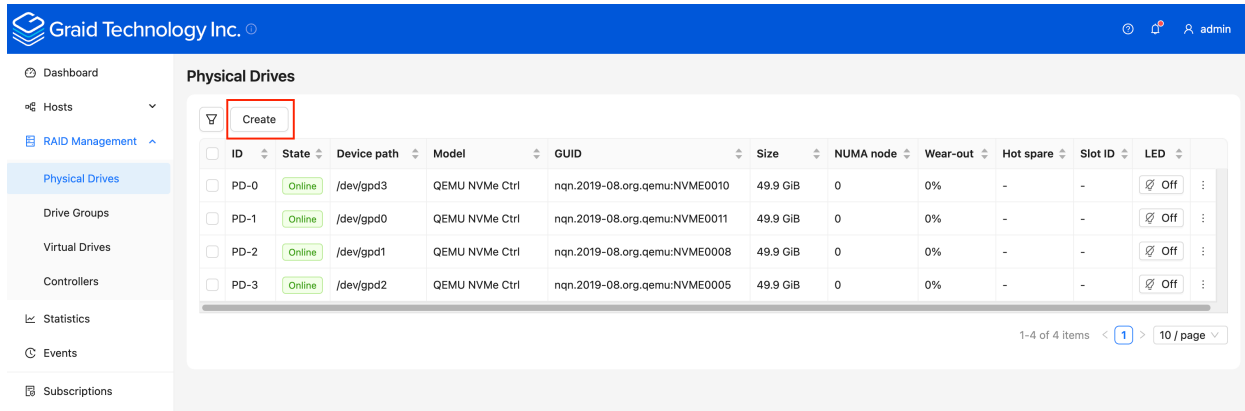
Step 7 Bind your GPU and start the RAID setup.



Managing Physical Drives

Creating a Physical Drive

To create physical drives, please log in to the SupremeRAID™ SE Management Console, then navigate to the RAID management / Physical Drives section on the sidebar menu and click the “Create” button.

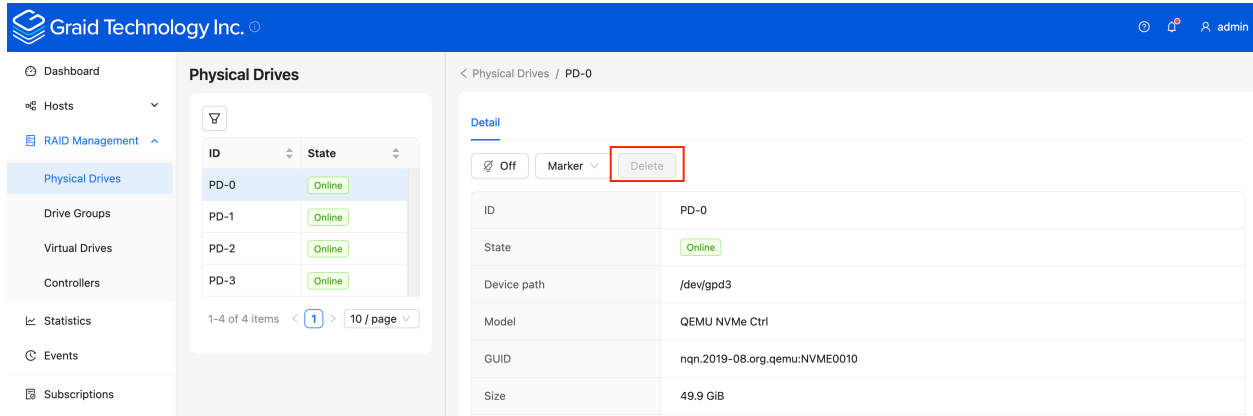


Physical Drive State:

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.
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.
COPYBACK	Physical Drive is performing copyback

Deleting a Physical Drive

To delete physical drives, please select the physical drives you want to delete and click the “Delete” button.

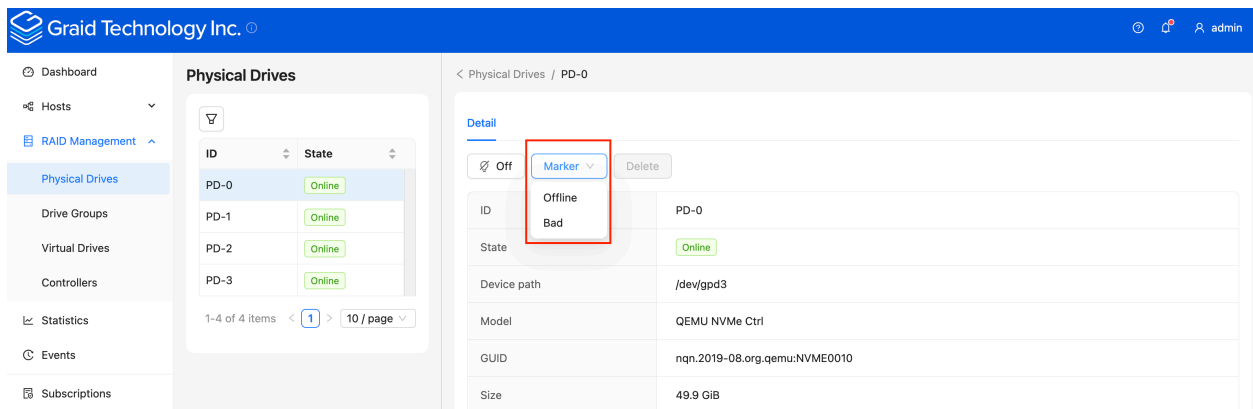


The screenshot shows the Graid Technology Inc. web interface. On the left is a navigation menu with options like Dashboard, Hosts, RAID Management, Physical Drives, Drive Groups, Virtual Drives, Controllers, Statistics, Events, and Subscriptions. The main content area is titled 'Physical Drives' and shows a list of four physical drives (PD-0, PD-1, PD-2, PD-3), all with a state of 'Online'. Below the list is a pagination control showing '1-4 of 4 items' and '10 / page'. To the right, the 'Detail' view for 'PD-0' is shown, with a 'Delete' button highlighted by a red box. The details for PD-0 include: ID: PD-0, State: Online, Device path: /dev/gpd3, Model: QEMU NVMe Ctrl, GUID: nqn.2019-08.org.qemu:NVME0010, and Size: 49.9 GiB.

Note: You cannot delete a physical drive which is used by drive group.

Marking a Physical Drive Online/Offline/Good/Bad

To mark a physical drive as Online/Offline/Good/Bad, please select the physical drives you want to change and click the “Marker” button.



The screenshot shows the Graid Technology Inc. web interface. On the left is a navigation menu with options like Dashboard, Hosts, RAID Management, Physical Drives, Drive Groups, Virtual Drives, Controllers, Statistics, Events, and Subscriptions. The main content area is titled 'Physical Drives' and shows a list of four physical drives (PD-0, PD-1, PD-2, PD-3), all with a state of 'Online'. Below the list is a pagination control showing '1-4 of 4 items' and '10 / page'. To the right, the 'Detail' view for 'PD-0' is shown, with a 'Marker' button highlighted by a red box. The details for PD-0 include: ID: PD-0, State: Online, Device path: /dev/gpd3, Model: QEMU NVMe Ctrl, GUID: nqn.2019-08.org.qemu:NVME0010, and Size: 49.9 GiB.

Assigning a Hot Spare Drive

To assign a physical drive as a hot spare, please select the physical drives you want assign and click the “Hot spare” button.

The screenshot displays the Graid Technology Inc. web interface. On the left is a navigation menu with options like Dashboard, Hosts, RAID Management, Physical Drives, Drive Groups, Virtual Drives, Controllers, Statistics, Events, and Subscriptions. The main content area is titled 'Physical Drives' and shows a table of physical drives:

ID	State
PD-0	Unconfigured good
PD-1	Online
PD-2	Online
PD-3	Online

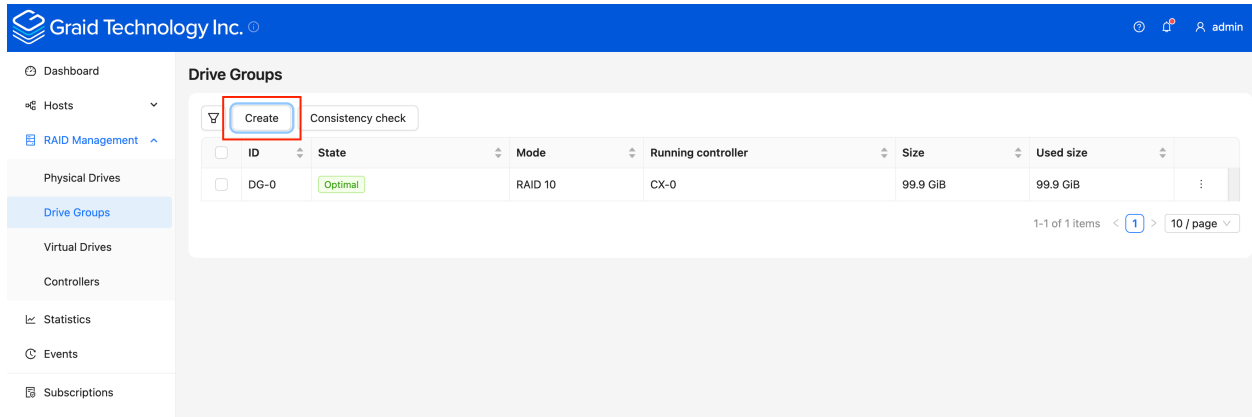
Below the table, it indicates '1-4 of 4 items' and '10 / page'. To the right, the 'Detail' view for 'PD-0' is shown, including buttons for 'On', 'Mark', 'Hot spare' (highlighted with a red box), and 'Delete'. The details table for PD-0 is as follows:

ID	PD-0
State	Unconfigured good
Device path	/dev/gpd3
Model	QEMU NVMe Ctrl
GUID	nqn.2019-08.org.qemu:NVME0010
Size	49.9 GiB

Managing Drive Groups

Creating a Drive Group

To create a drive group, please log in to the SupremeRAID™ SE Management Console, then navigate to the RAID management / Drive Group section on the sidebar menu and click the “Create” button.



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 (cp)	Drive group is performing copyback.
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.

State	Description
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.

Deleting a Drive Group

To delete a drive group, please log in to the SupremeRAID™ SE Management Console, then navigate to the RAID management / Drive Group section on the sidebar menu. Please select the drive group you want to delete and click the “Delete” button.

The screenshot shows the SupremeRAID™ SE Management Console interface. The left sidebar contains navigation options: Dashboard, Hosts, RAID Management (selected), Physical Drives, Drive Groups (selected), Virtual Drives, Controllers, Statistics, Events, and Subscriptions. The main content area is titled 'Drive Groups' and shows a list of drive groups. A table lists the drive group 'DG-0' with state 'Optimal' and mode 'RAID 10'. Below the table, there is a '1-1 of 1 items' and '10 / page' indicator. The 'Detail' tab is active, showing a 'Delete' button highlighted with a red box. The details for 'DG-0' are as follows:

ID	DG-0
State	Optimal
Mode	RAID 10
Size	99.9 GiB
Used size	99.9 GiB
Running controller	CX-0

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

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 from the drive group, force the drive group to go online using Rescue mode.

The screenshot displays the RAID Management interface. On the left, a sidebar menu includes Dashboard, Hosts, RAID Management, Physical Drives, Drive Groups, Virtual Drives, Controllers, Statistics, Events, and Subscriptions. The main area is titled 'Drive Groups' and contains a table with columns for ID, State, and Mode. The table lists drive groups DG-0 through DG-7. DG-1 is highlighted with a red box, showing a 'Failed' state and 'RAID5' mode. To the right, the 'Detail' view for DG-7 is shown, with a 'Rescue mode' button highlighted in blue. The details for DG-7 include: Drive group ID: DG-7, State: Failed, Mode: RAID5, Size: 6.83 GiB, Used size: 3.42 GiB, Running controller: CX-0, Strip size: 4 KiB, and Virtual drive number: 1.

ID	State	Mode
DG-0	Optimal	RAID5
DG-1	Failed	RAID5
DG-2	Partially degraded	RAID5
DG-3	Offline	RAID5
DG-4	Recovery	RAID5
DG-5	Init	RAID5
DG-6	Resync	RAID5
DG-7	Optimal	RAID5

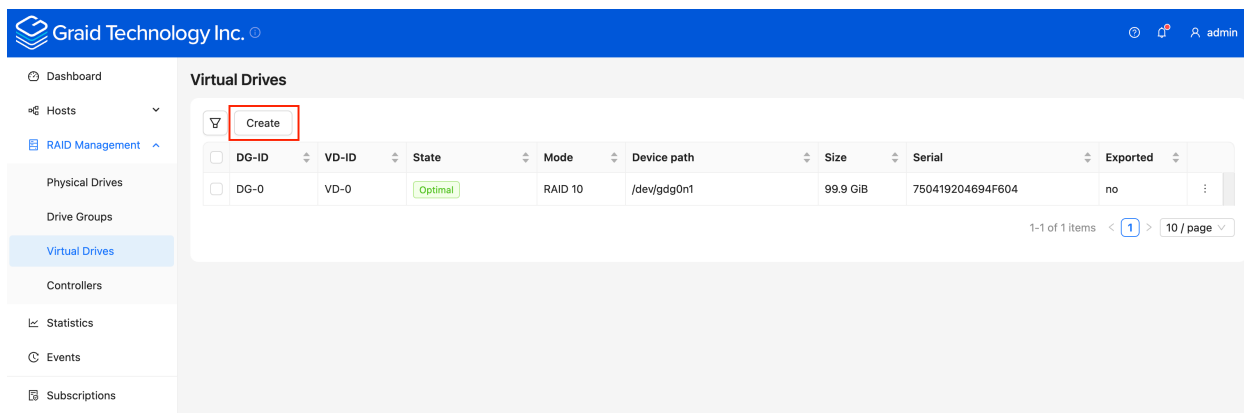
Property	Value
Drive group ID	DG-7
State	Failed
Mode	RAID5
Size	6.83 GiB
Used size	3.42 GiB
Running controller	CX-0
Strip size	4 KiB
Virtual drive number	1

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

Managing Virtual Drives

Creating a Virtual Drive

To create virtual drives, please log in to the SupremeRAID™ SE Management Console, then navigate to the RAID management / Virtual Drives section on the sidebar menu. Please click the “Create” button to select which the drive group you want to create virtual drives.



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.

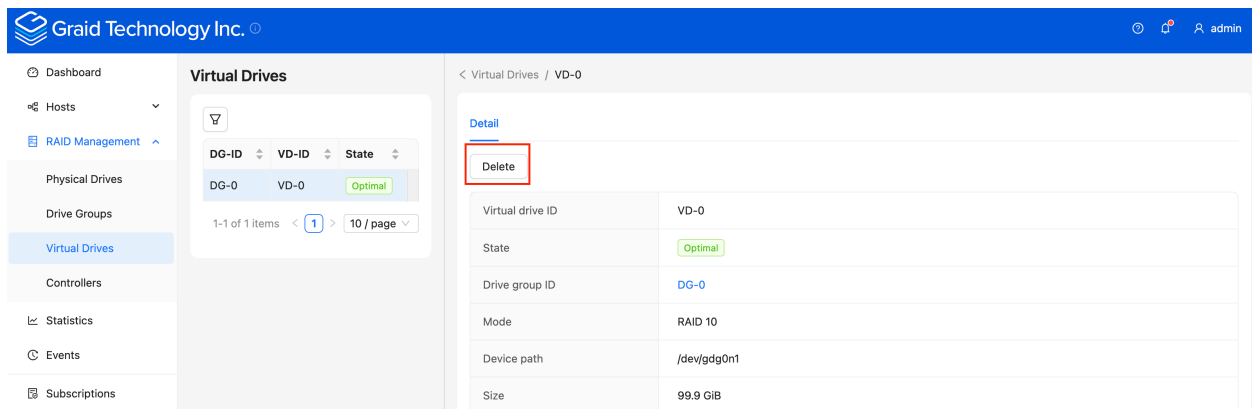
State	Description
RESCUE	Drive group is in rescue mode.

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 a Virtual Drive

To delete a virtual drive, please log in to the SupremeRAID™ SE Management Console, then navigate to the RAID management / Virtual Drives section on the sidebar menu. Please select the virtual drive you want to delete and click the “Delete” button.

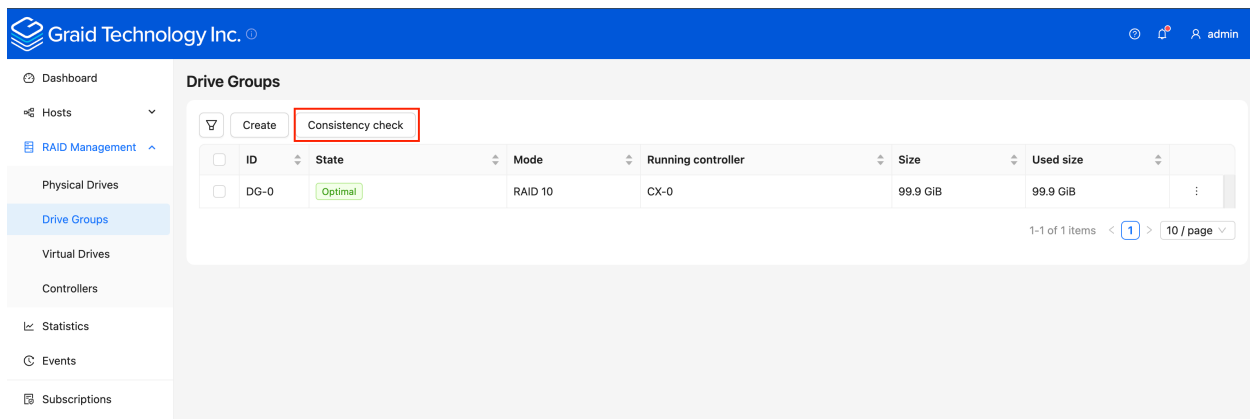


Features Overview

Consistency Checks

The consistency check operation verifies that the data is correct in Drive Groups that use RAID levels 1, 5, 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.

To start the consistency check, please log in to the SupremeRAID™ SE Management Console, then navigate to the RAID management / Drive Group section on the sidebar menu. Select the drive group you want to perform the Consistency Check on and click the “Consistency Check” button.



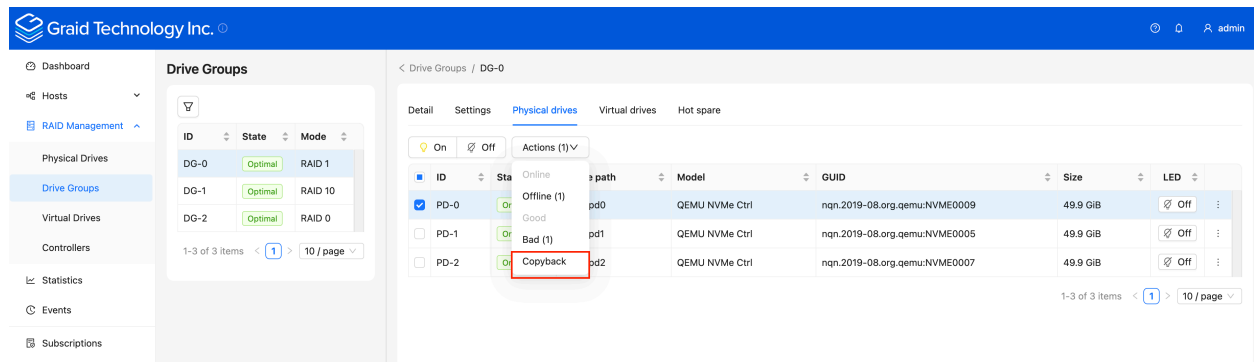
Note: You cannot perform a consistency check on RAID 0 because it does not provide data redundancy. Additionally, a consistency check can only run when the DG is in OPTIMAL or PARTIALLY_DEGRADED state.

Drive Copyback

The Drive Copyback feature allows users to manually initiate data migration from one drive to another without affecting the overall Drive Group state. This operation is user-controlled and can be performed for various reasons, such as replacing an aging drive, preparing for hardware upgrades, or managing storage configurations.

To start the Copyback, please log in to the SupremeRAID™ SE Management Console, then navigate to the RAID management / Drive Group section on the sidebar menu. Select the drive group you want to

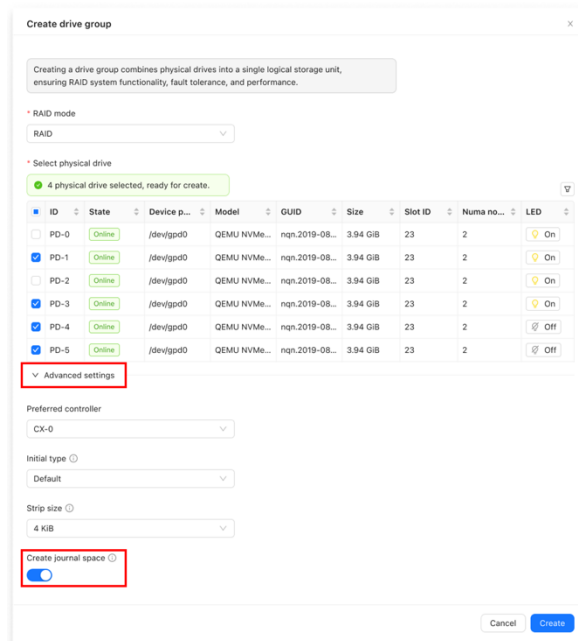
perform the Copyback on and click the Physical Drives tab. Choose the physical drive first and click the “Actions” button to conduct Copyback.



Double Failure Protection with Distributed Journaling

SupremeRAID™ SE incorporates a distributed journaling mechanism specifically designed to safeguard data during abnormal shutdowns in double-failure scenarios. This system ensures data integrity by logging data in a dedicated journaling space before writing it to the storage area. Any incomplete I/O operations are replayed upon service restart to maintain data consistency.

This journaling feature is automatically enabled in degraded mode to uphold data integrity. Additionally, users still have the flexibility to bypass journaling space reservation when creating a drive group.



Note: Only RAID5 can enable the journal function. If the user bypasses the creation of the journal space, it cannot be recreated.

Support for the Dataset Management (DSM) Deallocate command on virtual drives

The SupremeRAID™ SE driver introduces support for the NVMe DSM deallocate (trim) command on virtual drives, improving the efficiency of unused storage space management on NVMe SSDs. This feature allows filesystems or applications to issue deallocate commands on virtual drives, which are then translated by the driver and sent directly to the SSDs. By enabling the drives to manage deallocated blocks internally, this reduces write amplification, optimizes storage efficiency, and enhances overall performance.

When a discard command is issued to a virtual drive, it triggers a corresponding deallocate command to the underlying NVMe SSDs. The system supports a minimum discard range of 4KB, aligned with the logical block addressing (LBA) size, and can handle a maximum deallocate range of approximately 400 GiB per command. For larger discard operations, the filesystem and block layer handle the process seamlessly.

This feature is automatically enabled on NVMe SSDs that support the deallocate command and guarantee that deallocated blocks return zeros. For SSDs without this guarantee, the system defaults to a "write zeros" command to ensure data consistency. This flexible approach ensures broad compatibility across different SSDs while optimizing their individual capabilities.

To ensure the filesystem can take advantage of this capability and issue discard commands when files are deleted, it must be mounted with the **discard** option.

Attachments

Events for SupremeRAID™ SE

Category	Severity	Description
Physical Drive	Warning	Physical Drive <PD_ID> state has transitioned from <STATE_OLD> to unconfigured bad.
	Critical	Physical Drive <PD_ID> state has transitioned from <OLD_STATE> to failed.
	Warning	Physical Drive <PD_ID> state has transitioned from <OLD_STATE> to offline.
	Critical	Physical Drive <PD_ID> state has transitioned from <OLD_STATE> to missing.
	Info	Physical Drive <PD_ID> state has transitioned from <OLD_STATE> to online.
	Info	Physical Drive <PD_ID> state has transitioned from <OLD_STATE> to rebuild.
	Info	Physical Drive <PD_ID> state has transitioned from <OLD_STATE> to unconfigured good.
	Info	Physical Drive <PD_ID> has been successfully created.
	Info	Physical Drive <PD_ID> has been deleted.
	Info	Physical Drive <PD_ID> has been hot-plugged.
	Warning	Physical Drive <PD_ID> has been hot-removed.
	Warning	The temperature of Physical Drive <PD_ID> is currently <CURRENT_TEMP> degrees, which exceeds the Warning threshold of <THRESHOLD_TEMP> degrees. Critical Warning error code: ERROR_CODE.
	Critical	The temperature of Physical Drive <PD_ID> is currently <CURRENT_TEMP> degrees, which exceeds the Critical threshold of <THRESHOLD_TEMP> degrees. Critical Warning error code: ERROR_CODE.
	Critical	The available spare capacity <AVAIL_SPARE> of Physical Drive <PD_ID> has fallen below the threshold <SPARE_THRESHOLD>. Critical Warning error code: <ERROR_CODE>.
Critical	The NVM subsystem reliability of Physical Drive <PD_ID> has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability. Critical Warning error code: <ERROR_CODE>.	

	Critical	All of the media of Physical Drive <PD_ID> has been placed in read only mode. Critical Warning error code: <ERROR_CODE>.
	Critical	The volatile memory backup device of Physical Drive <PD_ID> has failed. Critical Warning error code: <ERROR_CODE>.
	Critical	The Persistent Memory Region of Physical Drive <PD_ID> has become read-only or unreliable. Critical Warning error code: <ERROR_CODE>.
	Warning	Physical Drive <PD_ID> is currently experiencing a wearout level of WEAROUT, surpassing the Warning threshold of <THRESHOLD_WEAROUT>.
	Critical	Physical Drive <PD_ID> is currently experiencing a wearout level of WEAROUT, surpassing the Critical threshold of <THRESHOLD_WEAROUT>.
Drive Group	Fatal	Drive Group <DG_ID> state has transitioned from <OLD_STATE> to failed.
	Critical	Drive Group <DG_ID> state has transitioned from <OLD_STATE> to offline.
	Critical	Drive Group <DG_ID> state has transitioned from <OLD_STATE> to degraded.
	Warning	Drive Group <DG_ID> state has transitioned from <OLD_STATE> to rescue.
	Warning	Drive Group <DG_ID> state has transitioned from <OLD_STATE> to partially degraded.
	Info	Drive Group <DG_ID> state has transitioned from <OLD_STATE> to optimal.
	Info	Drive Group <DG_ID> state has transitioned from <OLD_STATE> to recovery.
	Info	Drive Group <DG_ID> state has transitioned from <OLD_STATE> to init.
	Info	Drive Group <DG_ID> state has transitioned from <OLD_STATE> to resync.
	Info	Drive Group <DG_ID> has been successfully created.
	Info	Drive Group <DG_ID> has been deleted.
	Info	Consistency Check for Drive Group <DG_ID> has been manually aborted.
	Info	Consistency Check for Drive Group <DG_ID> has been aborted due to the deletion of the Drive Group.
	Info	Consistency Check for Drive Group <DG_ID> was aborted due to the Drive Group migrating from Controller <CX_OLD> to <CX_NEW>.
	Info	Consistency Check for Drive Group <DG_ID> has been aborted due to the Drive Group's state transitioning to <DG_STATE>.

	Info	Manual Consistency Check for Drive Group <DG_ID> has been completed.
	Info	Scheduled Consistency Check for Drive Group <DG_ID> has completed.
	Info	Manual Consistency Check for Drive Group <DG_ID> has started.
	Info	Scheduled Consistency Check for Drive Group <DG_ID> has started.
	Info	Inconsistency in Drive Group <DG_ID> has been fixed at: Drive Group block range: <DG_INTERS>.
	Critical	Inconsistency detected in Drive Group <DG_ID> at: Drive Group block range: <DG_INTERS>.
	Critical	Consistency Check for Drive Group <DG_ID> has been aborted due to the 'stop_on_error' policy.
	Critical	Consistency Check for Drive Group <DG_ID> has been aborted due to numerous inconsistencies found and fixed.
	Info	Journal Replay for Drive Group <DG_ID> has started.
	Info	Journal Replay for Drive Group <DG_ID> has been completed. Entry replayed <REPLAYNR>.
	Critical	Journal Replay for Drive Group <DG_ID> has been waiting Physical Drive <PD_ID> to be active.
	Critical	Journal Replay for Drive Group <DG_ID> has been aborted due to inconsistency detected on journal.
Virtual Drive	Info	Inconsistency for Virtual Drive <VD_ID> within Drive Group <DG_ID> has been fixed at: Virtual Drive block range: <VD_OFFSETS>.
	Critical	Inconsistency found in Virtual Drive VD_ID of Drive Group <DG_ID> at: Virtual Drive block range: <VD_OFFSETS>.
	Info	Virtual Drive VD_ID for Drive Group <DG_ID> has been created successfully.
	Info	Virtual Drive VD_ID for Drive Group <DG_ID> has been deleted.
	Info	Stripe cache for Virtual Drive <VD_ID> on Drive Group <DG_ID> has been deleted.
	Info	Stripe cache for Virtual Drive <VD_ID> on Drive Group <DG_ID> has been created successfully.
Controller	Warning	The temperature of Controller <CX_ID> is currently <CURRENT_TEMP> degrees, which exceeds the GPU threshold of <THRESHOLD_TEMP> degrees.

Warning	The temperature of Controller <CX_ID> is currently <CURRENT_TEMP> degrees, which exceeds the GPU memory threshold of <THRESHOLD_TEMP> degrees.
Warning	The temperature of Controller <CX_ID> is currently <CURRENT_TEMP> degrees, it will cause controller slowdown.
Critical	The temperature of Controller <CX_ID> is currently <CURRENT_TEMP> degrees, it will cause controller shutdown.