

Dell PowerEdge RAID Controller 12 User's Guide

PERC H965i Adapter, PERC H965i Front, PERC H965i MX, and PERC H965e Adapter

Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Chapter 1: Dell Technologies PowerEdge RAID Controller 12.....	9
Features of PERC H965i Adapter.....	10
Features of PERC H965i Front.....	10
Features of PERC H965i Front DC-MHS.....	11
Features of PERC H965e Adapter.....	13
Features of PERC H965i MX.....	14
Technical specifications of PERC 12 cards.....	15
Operating systems supported by PERC 12 cards.....	16
Chapter 2: Applications and User Interfaces supported by PERC 12.....	17
Comprehensive Embedded Management	17
Dell OpenManage Storage Management.....	17
Human Interface Infrastructure Configuration Utility.....	18
The PERC Command Line Interface.....	18
Chapter 3: Features of PowerEdge RAID Controller 12.....	19
Controller features.....	19
Hardware Root of Trust	19
Security Protocol and Data Model (SPDM).....	19
Device enumeration.....	19
UEFI Secure Boot.....	20
Auto-Configure Behavior (Execute Once)	21
Auto-Configure Behavior (Primary and Secondary settings)	21
Drive Performance Monitoring (DPM).....	21
Disk roaming.....	22
Hardware Accelerated I/O.....	22
Non-RAID disks.....	23
Physical disk power management.....	23
Firmware update.....	23
Snapdump.....	23
Physical disk coercion.....	24
PERC H965e support for PowerVault MD24XX enclosure.....	24
Virtual disk features.....	26
Auto-import foreign virtual drives.....	26
Virtual disk migration.....	26
Online capacity expansion.....	26
Virtual disk write cache and read-ahead policy.....	27
Virtual disk initialization.....	28
Full initialization.....	28
Fast initialization.....	28
Background operations.....	28
Background initialization.....	28
Consistency checks.....	29
Drive features.....	29

Self-Encrypting Drives (SEDs).....	29
Instant secure erase.....	30
Cryptographic Erase with PSID Revert.....	30
Physical disk erase.....	30
Sanitize Erase operation.....	30
4KB sector drives.....	31
Non-Volatile Memory Express.....	31
Fault tolerance.....	32
The SMART feature.....	32
Patrol Read.....	32
Physical disk failure detection.....	32
Controller cache.....	34
Battery (Energy Pack) Transparent Learn Cycle.....	35
Chapter 4: Install and remove a PERC 12 card.....	36
Before working inside your system.....	36
Safety instructions.....	36
Remove the PERC H965i Adapter.....	37
Install the PERC H965i Adapter.....	38
Remove the PERC H965i Front card.....	39
Install the PERC H965i Front card.....	40
Remove PERC H965i MX.....	42
Install PERC H965i MX.....	43
Removing the PERC H965e Adapter card.....	44
Installing the PERC H965e Adapter card.....	45
Chapter 5: Driver support for PERC H965i Adapter, PERC H965i Front, PERC H965i MX, and PERC H965e Adapter.....	47
Creating the device driver media.....	47
Download and save PERC H965i Adapter, PERC H965i Front, PERC H965i MX, and PERC H965e Adapter, drivers from the support site.....	47
Download and save PERC H965i Adapter, PERC H965i Front, PERC H965i MX, and PERC H965e Adapter drivers from the Dell Systems Service and Diagnostic Tools	48
Windows driver installation.....	48
Install PERC H965 series card driver during a Windows Server 2025 installation.....	48
Install PERC H965i Adapter, PERC H965i Front, PERC H965i MX, PERC H965e Adapter while newly installing Windows.....	49
Install PERC H965i Adapter, PERC H965i Front, PERC H965i MX, and PERC H965e Adapter driver on which Windows is already installed.....	49
Update PERC H965i Adapter, PERC H965i Front, PERC H965i MX, and PERC H965e Adapter driver that runs on Microsoft Windows operating system.....	49
Linux driver installation.....	50
Install or update a RPM driver package using the KMOD support.....	50
Install or update a RPM driver package using the KMP support.....	51
Load the driver while installing an operating system.....	51
Chapter 6: Manage the PERC 12 firmware.....	53
Upgrade firmware controller using Dell Update Package (DUP).....	53
Chapter 7: Manage PERC 12 controllers using HII Configuration Utility.....	54

Enter the PERC 12 HII Configuration Utility.....	54
Exit the PERC 12 HII Configuration Utility.....	54
Navigate to the PERC 12 Configuration Utility.....	55
View the HII Configuration Utility dashboard.....	55
Configuration management.....	55
Create virtual disk and configure virtual disk parameters.....	55
Select physical disks for creating VDs.....	56
Create profile based virtual disk.....	57
View disk group properties.....	57
View global hot-spare devices.....	58
Convert to Non-RAID disk.....	58
Delete configurations.....	58
Controller management.....	58
Save persistent events.....	58
Clear persistent events.....	58
Manage snapdump.....	59
Enable security.....	59
Disable security.....	59
Change security settings.....	60
Restore factory default settings.....	60
Manage the SAS Storage link speed.....	60
Manage the PCIe storage interface.....	60
Auto-Configure Behavior Management.....	61
Advanced controller properties.....	61
Virtual disk management.....	64
Virtual disk numbering.....	64
Configure Virtual Disks.....	66
Perform Online capacity expansion—Expand a VD within a disk group.....	66
Perform Online capacity expansion—Expand a VD by adding a PD to a disk group.....	66
Check consistency of VDs.....	67
Device management.....	67
View enclosure or backplane properties.....	67
View physical disk properties.....	68
Cryptographic erase.....	70
Cryptographic Erase with PSID Revert.....	70
Physical disk erase.....	71
Sanitize Block or Overwrite Erase operation.....	71
Assigning a global hot-spare.....	72
Assign a dedicated hot-spare.....	72
Convert a physical disk to a non-RAID disk.....	73
Hardware components.....	73
View battery (Energy Pack) properties.....	73
Security key management in HII configuration utility.....	74
Chapter 8: Security key and RAID management.....	75
Security key implementation.....	75
Local Key Management.....	75
Create a security key.....	75
Change security settings.....	76
Disable security key.....	76

Create a secured virtual disk.....	77
Secure a non-RAID disk.....	77
Secure a pre-existing virtual disk.....	77
Import a secured non-RAID disk.....	78
Import a secured virtual disk.....	78
Dell Technologies OpenManage Secure Enterprise Key Manager (SEKM).....	78
Supported controllers for OpenManage Secure Enterprise Key Manager (SEKM).....	79
Manage the Server Enterprise Key Manager (SEKM) feature.....	79
Disable the Secure Enterprise Key Manager (SEKM) feature.....	79
Manage virtual disks in the SEKM mode.....	79
Manage non-RAID disks in the SEKM mode.....	79
Transition of drives from Local Key Management to Secure Enterprise Key Management.....	79

Chapter 9: Troubleshooting issues in PERC 12 series cards..... 80

Virtual drive Blink and Unblink operations are unavailable in iDRAC UI and HII	80
Samsung PM893 and Samsung PM897 SATA SSDs are not supported on PERC12 firmware versions earlier than 8.8.0.0.18-26.....	81
The SMART data of SATA SSDs is unavailable from iDRAC	81
The default HDD cache policy is not updated to support the read-ahead operation.....	81
Firmware downgrade operation may fail if the firmware version downgraded to does not have the supported feature set.....	81
Unable to downgrade the firmware version of PERC12 to an older revision	82
When secured drives are hot-removed or hot-inserted in the server, iDRAC doesn't generate the PD Unlock event	82
During a Key Exchange Failure scenario in SEKM mode, native locked drives are Identified as "Locked with Foreign Key" in all interfaces.....	82
After importing the foreign configuration of the associated global hot-spare drive on a PERC12 series controller, the drive is still reported as both Online and global hot-spare	82
LED Operations using ESXCLI Commands for H965 and HBA465 Series controller is not available using VMware ESXi OS	83
Backplane slot count may be reported as 32 if there is an error with the backplane during discovery or runtime	83
Single virtual disk performance or latency in hypervisor configurations.....	83
Unable to discover or detect a PERC card.....	84
Unable to install the Microsoft Windows operating system	84
Only UEFI is supported and not BIOS.....	84
A fault firmware state is detected.....	84
Foreign configuration not found in HII.....	84
Degraded state of virtual disks.....	84
Boot-time errors.....	85
Event log errors.....	87
Application issues.....	96
Memory errors.....	100
Preserved Cache state.....	101
Security key errors.....	101
Secured foreign import errors.....	101
Failure to select or configure non Self-Encrypting Disks non-SED.....	101
Failure to delete security key.....	101
Failure of Cryptographic Erase on encryption-capable physical disks.....	101
General issues.....	102
PERC goes into Safe and Unresponsive state.....	102

Enclosure state is indicated as faulty instead of need attention.....	102
perccli2 commands and OS DUPs must have the controller host driver 8.0.0.69.0 for proper functioning.....	102
PERC card has yellow bang in Windows operating system device manager.....	102
PERC card not seen in operating systems.....	102
Unlocked drives if not imported appear as locked.....	102
Physical disk issues.....	103
Lifecycle log indicates drive number as unknown.....	103
An NVMe drive is not indicated as unsupported when EEDP is enabled.....	103
Unexpected generation of incorrect physical AEN.....	103
Unconfigured drives are not listed in the OS or Hardware Diagnostics.....	103
Physical disk in failed state.....	103
The PDR6 Error and Event message is saved in the log data when a new drive is inserted.....	103
Unable to rebuild a fault tolerant virtual disk.....	104
An NVMe drive is indicated as faulty when moving the drive to new backplane using hot-insertion..	104
Fatal error or data corruption reported.....	104
Multiple disks are inaccessible.....	104
Rebuilding data for a failed physical disk.....	104
Virtual disk fails during rebuild using a global hot-spare.....	105
Dedicated hot-spare disk fails during rebuild.....	105
Virtual disk fails rebuild using a dedicated hot spare.....	105
Physical disk takes a long time to rebuild.....	105
Drive removal and insertion in the same slot generates a foreign configuration event	105
SMART errors.....	106
SMART error detected on a non-RAID disk.....	106
SMART error detected on a physical disk in a non-redundant virtual disk.....	106
SMART error detected on a physical disk in a redundant virtual disk.....	106
Replace member errors.....	106
Source disk fails during replace member operation.....	107
Target disk fails during replace member operation.....	107
A member disk failure is reported in the virtual disk which undergoes replace member operation.....	107
Linux operating system errors.....	107
Virtual disk policy is assumed as write-through.....	107
Unable to register SCSI device.....	108
Drive indicator codes.....	108
The AutoSecure feature is not supported.....	109
HII error messages.....	109
Unhealthy Status of the drivers.....	109
Rebuilding a drive during full initialization.....	109
Chapter 10: Appendix—RAID description.....	110
Summary of RAID levels.....	110
RAID 10 configuration.....	111
RAID terminology.....	112
Disk striping.....	112
Disk mirroring.....	112
Spanned RAID levels.....	112
Parity data.....	113
Chapter 11: Getting help.....	114

Recycling or End-of-Life service information.....114
Contacting Dell.....114
Locating the Express Service Code and Service Tag.....114

Chapter 12: Documentation resources..... 116

Dell Technologies PowerEdge RAID Controller 12

Dell Technologies PowerEdge RAID Controller 12 or PERC 12 is a series of RAID disk array controllers that are developed by Dell for its PowerEdge servers. The PERC H965i Adapter, PERC H965i Front, PERC H965i MX, and PERC H965e Adapter controllers have the following characteristics:

- Reliability, high performance, and fault-tolerant disk subsystem management.
- RAID control capabilities including support for RAID levels 0, 1, 5, 6, 10, 50, and 60.
- Gen 4 PCIe x16 host interfaces.
- Supports Dell-qualified Serial Attached SCSI (SAS), Serial Advanced Technology Attachment (SATA), and PCIe SSD (NVMe) drives.
- Supports drive speeds for NVMe drives are 8 GT/s (Gen 3) and 16 GT/s (Gen 4) at maximum x2 lane width.
- Supports data rate throughput of 6 Gbps for SAS 2.0, 12 Gbps for SAS 3.0, and 22.5 Gbps for SAS 4.0 drives.
- Supports data rate of throughput of 6 Gbps for SATA 3.0 drives.

NOTE: The PERC H965i Front controller card is available in two types:

- PERC H965i Front controller card—Supported on 16th Generation (16G) PowerEdge servers.
- PERC H965i Front DC-MHS controller card—Supported on 17th (17G) Generation PowerEdge servers. However, both the types have separate Update Packages (DUPs).

NOTE:

- The PERC 12 controller tools such as PERCCLI2, drivers, and firmware are not backward compatible with previous versions of PERC controllers.
- Mixing disks of different speed (7,200 RPM, 10,000 RPM, or 15,000 RPM) and bandwidth (6 Gbps, 12 Gbps, or 24 Gbps) while maintaining the same drive type (SAS or SATA) and technology (HDD or SSD) is supported.
- PERC 12 controllers support only single SCSI LUN and single NVMe namespace devices. Multi-LUN and Multi-Namespace devices are not supported.
- PERC12 series controllers do not support BIOS, and are not listed in the BIOS Setup section. The Storage boot operation is available only in the UEFI mode.
- Shingled Magnetic Recording (SMR) drives are not supported on PERC 12 controllers.
- Mixing NVMe drives with SAS and SATA is not supported. Also, mixing HDD and SSD in a virtual disk is not supported.
- SAS4 "22.5 Gbps" speed is used synonymously with "24G" and "24 Gbps" in documents and some applications. 22.5 Gbps is the data rate and 24 Gbps is the link speed.
- For the safety, regulatory, and ergonomic information that is associated with these devices, and for more information about the Integrated Dell Remote Access Controller (iDRAC) or Lifecycle Controller remote management, see the platform-specific technical documentation.

CAUTION: For Japan Only—This is a Class A equipment. Operation of this equipment in a residential environment could cause Radio-wave interference. In this case, the user may be required to take corrective actions. (VCCI-A)

Topics:

- [Features of PERC H965i Adapter](#)
- [Features of PERC H965i Front](#)
- [Features of PERC H965i Front DC-MHS](#)
- [Features of PERC H965e Adapter](#)
- [Features of PERC H965i MX](#)
- [Technical specifications of PERC 12 cards](#)
- [Operating systems supported by PERC 12 cards](#)

Features of PERC H965i Adapter

Features of PERC H965i Adapter




Figure 1. Features of PERC H965i Adapter

- 1. Heatsink
- 2. Backplane Connector-B
- 3. Battery (Energy Pack)
- 4. Backplane Connector-A
- 5. PCIe connector

The following image shows another type of H965i adapter that you can use in a PowerEdge sled model. For example, PowerEdge C6620. In this type of adapter, you can detach the battery and install it in the air shroud of the sled. For information about installing battery into the air shroud, see [Install the PERC H965i Adapter](#).



Figure 2. PERC H965i Adapter with a detachable battery

Features of PERC H965i Front

Shows an image of the PERC H965i Front controller card. The PERC H965i Front controller card is available in two types. 1) The PERC H965i Front controller card is supported on 16th Generation (16G) PowerEdge servers. 2) The PERC H965i Front

DC-MHS controller card is supported on 17th (17G) Generation PowerEdge servers. Both controller types have separate Update Packages (DUPs). However, the procedure for installation and removal of both the controller cards is same.




Figure 3. PERC H965i Front

- | | |
|------------------------------|--------------------------|
| 1. PCIe connector B | 2. PCIe connector A |
| 3. Backplane connector B | 4. Backplane connector A |
| 5. Battery (Energy Pack) | 6. Heat sink |
| 7. Power card edge connector | |

Features of PERC H965i Front DC-MHS

Shows an image of the PERC H965i Front DC-MHS controller card. The PERC H965i Front controller card is available in two types. 1) The PERC H965i Front controller card is supported on 16th Generation (16G) PowerEdge servers. 2) The PERC H965i Front DC-MHS controller card is supported on 17th (17G) Generation PowerEdge servers. Both controller types have separate Update Packages (DUPs). However, the procedure for installation and removal of both the controller cards is same.




Figure 4. PERC H965i Front DC-MHS

- | | |
|-----------------------------|--------------------------|
| 1. Heatsink | 2. Backplane connector B |
| 3. Battery (Energy Pack) | 4. Backplane connector A |
| 5. Powercard Edge connector | 6. PCIe connector B |
| 7. PCIe connector A | |

Features of PERC H965e Adapter

PERC H965e Adapter




Figure 5. PERC H965e Adapter

- 1. Fan cable connector
- 2. Heat sink
- 3. Fan
- 4. Battery
- 5. Battery cable connector
- 6. Port 2
- 7. Port 2
- 8. Port 1
- 9. Port 0

Features of PERC H965i MX




Figure 6. Features of PERC H965i MX

- 1. Battery (Energy Pack) under cover
- 2. Heatsink
- 3. Backplane connector A
- 4. Backplane connector B
- 5. PCIe cable connector

Technical specifications of PERC 12 cards

NOTE: The PERC H965i Front controller card is available in two types:

- PERC H965i Front controller card—Supported on 16th Generation (16G) PowerEdge servers.
- PERC H965i Front DC-MHS controller card—Supported on 17th (17G) Generation PowerEdge servers. However, both the types have separate Update Packages (DUPs).

The following table lists the specifications of PERC 12 cards.

Table 1. Technical specifications of PERC 12 cards

Feature	PERC H965i Adapter	PERC H965i Front	PERC H965i MX	PERC H965e Adapter
RAID levels	0, 1, 5, 6, 10, 50, 60	0, 1, 5, 6, 10, 50, 60	0, 1, 5, 6, 10, 50, 60	0, 1, 5, 6, 10, 50, 60
Non-RAID	Yes	Yes	Yes	Yes
Host bus type	16-lane, PCIe 4 at 16 Gbps	16-lane, PCIe 4 at 16Gbps	16-lane, PCIe 4 at 16 Gbps	16-lane, PCIe 4 at 16Gbps
Side-band Management	I2C, PCIe VDM	I2C, PCIe VDM	I2C, PCIe VDM	I2C, PCIe VDM
Enclosures per port	Not applicable	Not applicable	Not applicable	4
Processor	Broadcom RAID-on-chip, SAS4116W chipset			
Battery (Energy Pack) backup unit	Yes	Yes	Yes	Yes
Local Key Management security	Yes	Yes	Yes	Yes
Secure enterprise key manager security	Yes	Yes	Yes	Yes
Controller queue depth	8,192	8,192	8,192	8,192
Nonvolatile cache	Yes	Yes	Yes	Yes
Cache memory	8 GB DDR4 3200 MT/s cache			
Cache function	Write-back, read-ahead, write-through, always write-back, no read-ahead			
Max no of VD's in RAID mode	240	240	240	240
Max no of disk groups	64	64	64	64
Max number of VD's per disk group	16	16	16	16
Hot-swap devices supported	Yes	Yes	Yes	Yes
Auto-Configure behavior (Primary and Execute once)	Yes	Yes	Yes	Yes
Hardware XOR engine	Yes	Yes	Yes	Yes
Online capacity expansion	Yes	Yes	Yes	Yes

Table 1. Technical specifications of PERC 12 cards (continued)

Feature	PERC H965i Adapter	PERC H965i Front	PERC H965i MX	PERC H965e Adapter
Dedicated and global hot spare	Yes	Yes	Yes	Yes
Supported Drive Types	22.5 Gbps SAS, 12 Gbps SAS, and 6 Gbps SATA/SAS. Gen3 (8 GT/s) and Gen4 (16 GT/s) NVMe	22.5 Gbps SAS, 12 Gbps SAS, and 6 Gbps SATA/SAS. Gen3 (8 GT/s) and Gen4 (16 GT/s) NVMe	22.5 Gbps SAS, 12 Gbps SAS, and 6 Gbps SATA/SAS. Gen3 (8 GT/s) and Gen4 (16 GT/s) NVMe	22.5 Gbps SAS, 12 Gbps SAS
VD strip element size	Both 64 KB and 256 KB allowed on an HDD. Only 64 KB allowed on an SSD.	Both 64 KB and 256 KB allowed on an HDD. Only 64 KB allowed on an SSD.	Both 64 KB and 256 KB allowed on an HDD. Only 64 KB allowed on an SSD.	Both 64 KB and 256 KB allowed on an HDD. Only 64 KB allowed on an SSD.
NVMe PCIe support	Gen 4	Gen 4	Gen 4	Gen 4
Configuration maximum SAS/SATA physical disks support	<ul style="list-style-type: none"> Without SAS Expander: 16 drives per controller With SAS Expander: Limited by platform offerings 	<ul style="list-style-type: none"> Without SAS Expander: 16 drives per controller With SAS Expander: Limited by platform offerings 	Limited by platform: 8 drives per controller	240 SAS drives
Configuration maximum NVMe physical drives support	<ul style="list-style-type: none"> Without PCIe Switch Expander: 8 drives per controller With PCIe Switch Expander: Limited by platform offerings 	<ul style="list-style-type: none"> Without PCIe Switch Expander: Eight drives per controller With PCIe Switch Expander: Limited by platform offerings 	Limited by platform: Eight drives per controller	Not supported
Controller maximum physical disks supported	<ul style="list-style-type: none"> 240 SAS/SATA 24 NVMe 	<ul style="list-style-type: none"> 240 SAS/SATA 24 NVMe 	<ul style="list-style-type: none"> 240 SAS/SATA 24 NVMe 	240 SAS
Drive sector size supported	512B, 512e, and 4Kn	512B, 512e, and 4Kn	512B, 512e, and 4Kn	512B, 512e, and 4Kn
Storage Boot Support	UEFI-only	UEFI-only	UEFI-only	No boot support

Operating systems supported by PERC 12 cards

NOTE: A PERC H965e controller is not supported as a boot controller.

- For a list of supported operating systems by a specific server for the PERC 12 cards, see [Dell Technologies Enterprise operating systems support](#).
- For the latest list of supported operating systems and driver installation instructions, see the operating system documentation at [Operating System Documentation](#).
- For specific operating system service pack requirements, see the Drivers and Downloads section at [Dell Technologies support site](#).

Applications and User Interfaces supported by PERC 12

PERC 12 card management applications include the Comprehensive Embedded Management (CEM), Dell OpenManage Storage Management, the Human Interface Infrastructure (HII) Configuration Utility, and the PERC Command Line Interface (CLI). They enable you to manage and configure the RAID system, create and manage multiple disk groups, control and monitor multiple RAID systems, and provide online maintenance.

Topics:

- [Comprehensive Embedded Management](#)
- [Dell OpenManage Storage Management](#)
- [Human Interface Infrastructure Configuration Utility](#)
- [The PERC Command Line Interface](#)

Comprehensive Embedded Management

Comprehensive Embedded Management (CEM) is a storage management solution for Dell systems that enables you to monitor the RAID and network controllers that are installed on the system using iDRAC without an operating system installed on the system.

Using CEM enables you to do the following:

- Monitor devices with and without an operating system that is installed on the server.
- Provide a specific location to access monitored data of the storage devices and network cards.
- Allows controller configuration for all 12th generation of PERC 12 controllers.

i **NOTE:** If you boot the system to HII (F2) or Lifecycle Controller (F10), then you cannot view the PERC cards on the CEM UI. The PERC cards are displayed on the CEM UI only after the system boot is complete.

i **NOTE:** It is not recommended that you create more than eight VDs simultaneously using CEM.

i **NOTE:** The new features that are provided in a PERC12 controller version 8.11.0.0.18-22 can be used only if you have iDRAC 7.10.50.00 for 16G PowerEdge servers and 1.20.25.00 for 17G and later PowerEdge servers.

Dell OpenManage Storage Management

Dell OpenManage Storage Management is a storage management application for Dell servers that provides enhanced features for configuring locally attached RAID disk storage. It enables you to perform controller and enclosure functions for all supported RAID controllers and enclosures from a single User Interface (UI) or Command Line Interface (CLI).

The UI is wizard-driven with features for novice and advanced users, and provides Online Help (OLH). Using the Dell OpenManage Storage Management application you can protect your data by configuring data-redundancy, assigning hot-spares, or rebuilding failed physical disks. The fully featured CLI, which is available on selected operating systems, enables you to perform RAID management tasks either directly from the console or by running commands.

i **NOTE:** For more information, see the *Dell OpenManage Storage Management User's Guide* available at [OpenManage Manuals](#).

Human Interface Infrastructure Configuration Utility

The Human Interface Infrastructure (HII) configuration utility is a storage management application integrated into the system BIOS <F2>. It is used to configure and manage your Dell PowerEdge RAID Controller (PERC) virtual disks and physical disks. This utility is independent of the operating system.

The PERC Command Line Interface

The PERC Command Line Interface (CLI) is a storage management application. This utility allows you to set up, configure, and manage your Dell PowerEdge RAID Controller (PERC) by using the CLI.

 **NOTE:** For more information, see the *Dell PowerEdge RAID Controller CLI Reference Guide* and the *Dell PowerEdge RAID Controller 12 User's Guide* available on the support site.

Features of PowerEdge RAID Controller 12

Topics:

- Controller features
- Virtual disk features
- Virtual disk initialization
- Background operations
- Drive features
- Fault tolerance

Controller features

This section lists the following controller features supported on Dell Technologies PowerEdge RAID Controller 12 cards:

- Hardware Root of Trust
- Security Protocol and Data Model (SPDM)
- Device enumeration
- UEFI Secure Boot
- Auto-Configure Behavior (Execute Once)
- Auto-Configure Behavior (Primary and Secondary settings)
- Disk roaming
- Hardware Accelerated I/O
- Non-RAID disks
- Physical disk power management
- Firmware update
- Snapdump
- Physical disk coercion
- PERC H965e support for PowerVault MD24XX enclosure

Hardware Root of Trust

Hardware Root of Trust (RoT) builds a chain of trust by authenticating all the firmware components prior to its execution and permits only the authenticated firmware to be installed and upgraded. The controller boots from an Internal Boot ROM (IBR) that establishes the initial RoT and this process authenticates and builds a chain of trust with succeeding software using this RoT.

Security Protocol and Data Model (SPDM)

Security Protocol and Data Model (SPDM) is a mechanism by which iDRAC can verify the authenticity of the PERC in the system. Each PERC card is manufactured with a unique Device Identity certificate signed by Dell to ensure that the PERC is a Dell Certified controller. iDRAC will automatically retrieve a Device Identity certificate from the PERC card during boot and verify its identity against the PERC and notify the user if the device could not be authenticated. The SPDM feature is supported in the PERC 8.8.0.0-18-26 or later versions.

Device enumeration

All devices attached to the controller are assigned an ID from a persistent range of numbers. This includes backplanes, physical disks, and virtual disks. When a backplane is discovered, it is assigned a range of IDs based on the number of slots the backplane

has. Each slot is assigned a dedicated ID in ascending order from the reserved range. When a disk is inserted, an ID allocated to the slot will be assigned.

Virtual disks are assigned an ID based on creation. The first virtual disk created will be assigned ID 1 and increases for each virtual disk. If a virtual disk is deleted, a newly created virtual disk will reuse that ID before using the next free ID. The newly created virtual disk will only reuse that ID if more than 120 seconds have passed since the previous virtual disk was deleted.

- Operating system device enumeration
 - All devices are presented to the operating system as a SCSI device.
 - Virtual disks and non-RAID disks are presented to the OS in the order of their ID.
 - In case of some automated OS installation processes, by default, the OS is installed on the disk that is first presented to the system. If you are installing to a virtual disk, it is recommended to make the first created VD the OS. If you installing to a non-RAID disk, it is recommended to install the disk in slot 0 of the front backplane.
 - Only Virtual Disks and non-RAID disks are exposed to the OS. Unconfigured disks are hidden from the OS.

i NOTE: For Linux servers only:

- Using the Device Reporting Order feature, you can select the order in which you can present either VDs or non-RAID drives to the OS.
- Using the First Device feature, you can select the order in which the physical device must be presented to the OS. The physical device will be enumerated first during the OS installation (when the OS has the latest Inbox drivers).

i NOTE: Operating system enumeration may not be in this order if virtual disks or non-RAID disks are created while the operating system is running. The operating system may name devices based on the order in which they were created resulting in the operating system enumeration changing after reboot. It is recommended to restart the system for the final device enumeration after creating any virtual disk or non-RAID disk.

UEFI Secure Boot

UEFI Secure Boot is a technology that eliminates a major security void that may occur during a handoff between the UEFI firmware and UEFI operating system (OS). In UEFI Secure Boot, each component in the chain is validated and authorized against a specific certificate before it is allowed to load or run. Secure Boot removes the threat and provides software identity checking at every step of the boot-Platform firmware, Option Cards, and OS BootLoader. For more information about using UEFI Secure Boot in PowerEdge servers, see the "Secure Boot Configuration from BIOS Settings or F2" section of the iDRAC User's Guide available on the support site.

To view a list of supported operating systems, see [PowerEdge server supported operating systems](#). By default, Windows and VMware support UEFI Secure Boot on both in-box and out-of-box drivers. However, Linux supports Secure Boot only with in-box drivers. For Out-of-Box Linux drivers, you must install Broadcom public key in the UEFI or OS signature database. For information about installing signature database, see the technical documentation of the respective platform or OS. The Broadcom Public database key is provided with a OS driver package on the Dell support site. The following table lists the OS and their supported driver types and also indicates whether or not the UEFI Secure Boot feature is supported with PERC12.

Table 2. UEFI Secure Boot Supported Operating systems on PowerEdge servers using PERC 12

Operating System	Driver Type	Secure Boot Supported	Signature Database
Microsoft Windows Server	In-Box	Yes	Native
Microsoft Windows Server	Out-of-Box	Yes	Native
RHEL	In-Box	Yes	Native
RHEL	Out-of-Box	Yes (with DB Install)	Broadcom Public
SUSE SLES	In-Box	Yes	Native
SUSE SLES	Out-of-Box	Yes (with DB Install)	Broadcom Public
VMware	VMware-provided	Yes	Native
Ubuntu	In-Box	Yes	Native

Auto-Configure Behavior (Execute Once)

The Auto-Configure Behavior (execute once) operation configures eligible ready-state drives based on the selected configuration. The default Auto-Config Behavior is Off. The supported configuration options for this feature are:

- Single drive RAID 0 with write-through cache policy
- Single drive RAID 0 with write-back cache policy
- Non-RAID disk
- Secured single drive RAID 0 with write-through cache policy
- Secured single drive RAID 0 with write-back cache policy
- Secured non-RAID disk

NOTE: The secured configuration options are available only when controller security is enabled and SEDs are present.

NOTE: The cache policy settings do not apply to non-RAID drives.

Auto-Configure Behavior (Primary and Secondary settings)

The Auto-Configure Behavior feature is used to configure new unconfigured drives (ready-state drives) during reboot and hot-insertion based on the selected configuration options (primary and secondary settings). The primary setting of Auto-Configure Behavior is used until the maximum configuration for the selected option is reached. Secondary setting of Auto-Configure Behavior is used after the maximum configuration for primary Auto-Configure option is reached.

Primary and secondary Auto-Configure settings do not impact known unconfigured drives during boot and hot-insertion. Known unconfigured drives are drives that were previously configured as virtual disk, hot-spare, or non-RAID disk and are now in the Ready state because the configuration on those drives was deleted.

When primary and secondary Auto-Configure options are changed from Off to any other settings, all unconfigured drives (including known devices) present on the controller remain unconfigured or in ready state.

NOTE: Secured configuration options are available only when the controller security is enabled.

The following table provides options for the supported primary Auto-Configure Behavior:

Table 3. Auto-Configure Behavior settings

Settings	Description
Off	The Auto-Configure Behavior feature is disabled. All new unconfigured drives remain in unconfigured or ready state. Existing configured drives are not affected when primary, secondary, or both Auto-Configure settings are changed to Off . NOTE: Applications use Off and Ready terms interchangeably for this setting.
Non-RAID	New unconfigured drives are configured as non-RAID disks during boot or during hot-insertion. All known unconfigured drives and existing configured drives remain unaffected during boot, hot-removal, or reinsertion.
Secured Non-RAID Disk	New unconfigured drives are configured as non-RAID disks during boot or hot-insertion. All known unconfigured drives and existing configured drives remain unaffected during boot, hot-removal, or reinsertion. SEDs are secured with the controller security key.

NOTE: Secondary Auto-Configure settings may not be supported on certain drive types. In such cases, the only available secondary Auto-Configure behavior option is Off.

Drive Performance Monitoring (DPM)

Enables you to collect performance data about one or more drives associated with the controller. You can use the metric to identify one or more slow performing drives in the system. For more information about the DPM feature, see the *Dell PERC CLI Reference Guide* at [Storage Controller Documentation](#).

Disk roaming

Disk roaming is when a physical disk is moved from one cable connection or backplane slot to another on the same controller. The controller automatically recognizes the relocated physical disks and logically assigns them to the virtual disks that are part of the disk group. If the physical disk is configured as a non-RAID disk, then the relocated physical disk is recognized as a non-RAID disk by the controller.

CAUTION: It is recommended that you perform a disk-roaming operation only when the server is powered off. Do not move drives between slots when the server is running. If it is attempted, maintain a 4-second delay between the removal and insertion operation for SAS or SATA devices. Maintain an 8-second delay between the removal and insertion operation for NVMe devices. Inability to observe this delay may result in issues while detecting devices.

Using disk roaming

About this task

Perform the following steps to use disk roaming:

Steps

1. Power off the server, PDs, enclosures, and server components.
2. Disconnect power cables from the server.
3. Move the physical disks to desired positions on the backplane or the enclosure.
4. Perform a safety check. Make sure the physical disks are inserted properly.
5. Power on the server.

Results

The controller detects the RAID configuration from the configuration data on the physical disks.

Hardware Accelerated I/O

The PERC12 provides improved hardware accelerated I/O when compared to the previous generation of controllers. The following table lists the different types of I/O provided by PERC 12 and previous generation of storage controllers:

Table 4. Comparison of performance for different I/O type for PERC 12 and previous generation of controllers

I/O Type	VD Cache	Hardware Accelerated in previous generations	Hardware Accelerated in PERC12 generation
SAS/SATA non-RAID Read-Write	N/A	Yes	Yes
NVMe non-RAID Read Write	N/A	Yes	Yes
Any RAID Level Single-Strip-Read	Write-Through	Yes	Yes
R0, R1, R10 Single-Strip-Write	Write-Through	Yes	Yes
R0, R1, R10 Single-Strip-Read	Write-Back with cache miss	Yes	Yes
R0, R1, R10 Single-Strip-Write	Write-Back with cache miss	Yes	Yes
R5, R6 Single-Strip-Read	Write-Back with cache miss	Yes	Yes
Any RAID Level MultiStrip-Read	Write-Through	No	Yes
Any RAID Level MultiStrip-Read	Write-Back	No	Yes

Table 4. Comparison of performance for different I/O type for PERC 12 and previous generation of controllers (continued)

I/O Type	VD Cache	Hardware Accelerated in previous generations	Hardware Accelerated in PERC12 generation
R0, R1, R10 MultiStrip-Write	Write-Through	No	Yes
R0, R1, R10 Single-Strip-Read	Write-Back with cache hit	No	Yes
R0, R1, R10 Single-Strip-Write	Write-Back with cache hit	No	Yes
R5, R6 Write	Write-Through	No	Yes
R5, R6 Write	Write-Back	No	Yes
R5, R6 Single-Strip-Read	Write-Back with cache hit	No	Yes
Cache Flush	Write-Back	No	Yes
Degraded VD-Read	Write-Through and Write-Back	No	Yes

Non-RAID disks

A non-RAID disk is a single disk to the host, and not a RAID volume. The only supported cache policy for non-RAID disks is Write-Through, and cannot be altered. Non-RAID disks provide passthrough support allowing the operating system nearly full access to the drive. All drives are presented to the host as a SAS target. The controller will convert the SAS/SCSI command to the appropriate SATA or NVMe command.

Physical disk power management

Physical disk power management is a power-saving feature of PERC 12 series cards. The feature allows disks to be spun down based on disk configuration and I/O activity. The feature is supported on all rotating SAS and SATA disks, and includes unconfigured and hot-spare disks. By default, the physical disk power management feature is enabled on unconfigured drives and disabled on hot-spare drives.

Firmware update

Upgrade or downgrade the PERC 12 firmware using Update Packages or PERC CLI. See [Manage the PERC 12 firmware](#).

Secure firmware update

Enables you to upgrade firmware by using an RSA encryption-decryption algorithm. The firmware upgrade operations you perform on a PERC 12 controller are secure. You can upgrade the PERC firmware by using only a Dell certified firmware.

Rebootless firmware update

Starting from PERC 12 generation of controllers, rebootless firmware update is supported in most cases. However, a system-reset operation may still be required to activate the newly downloaded firmware in exceptional cases such as presence of preserved cache or update to firmware components related to power handling or caching.

Snapdump

The Snapdump feature provides the Dell Technical Support team with the debug information that can help to find the cause of firmware failure. In the instance of firmware failures, the firmware collects the logs and information at the time of failure, which are stored in a compressed file called a snapdump.

Snapdumps are also generated manually to provide additional debug information. When a snapdump is generated, it is stored in the controller's cache memory. In the event of a power loss the controller will offload the snapdump as part of its cache preservation mechanism. Snapdumps are preserved by default through four restarts before it is deleted.

NOTE: The I/O operation will pause for about 5-10 seconds while collecting the on-demand snapdump.

To generate, delete, or download a stored snapdump, see the *Dell PowerEdge RAID Controller CLI Reference Guide* available on the support site.

Physical disk coercion

Enables you to create consistent disk size for drives that have slightly different sizes. Use the coercion feature in HII to match the sizes of drives that are slightly larger with the smaller drives in case you want to insert slightly smaller drives later. PERC 12 controllers support coercion modes such as **None (32 MiB)**, **128 MiB**, or **1 GiB**. Default mode is **128 MiB**. For virtual disks and hot-spares, the last **512 MiB** is used for configuration information.

Coerced size = (raw disk size minus 512 MiB configuration information) rounded down to the nearest full coercion mode size bucket.

Set Drive Coercion Mode

To provide consistent disk size for drives that are slightly different in sizes, set the drive coercion mode in HII.

Steps

1. Log in to HII.
2. Go to **Main Menu > Controller Management > Advanced Controller Properties**.
3. Select the **Drive Coercion Mode** and click **Apply Changes**.

PERC H965e support for PowerVault MD24XX enclosure

The PERC H965e adapter card is an external card and supports MD24xx series disk-attached external storage enclosure. The PERC H965e adapter supports up to eight MD2412 or MD2424 enclosures in a multipath configuration and up to four MD2460 enclosures in a multipath configuration. The PERC H965e adapter supports only symmetric enclosure topology. To set up a configuration with redundant paths—port 0 and 1, or port 2 and 3—a controller must be cabled to the ports of a single enclosure.

NOTE: The H965e adapter card does not support single path configuration.

NOTE: The H965e card supports redundant paths when used with Dell PowerVault MD2412, MD2424, and MD2460 series disk-attached storage enclosures.

NOTE: For PowerVault MD24XX series disk-attached storage enclosure, the multipath configuration is the only valid configuration.

Connect an H965e adapter to MD24xx enclosures

To set up a multipath enclosure on the PERC H965e card:

- On the controller, port 0 and 2, or port 1 and 3 must be used together. Interleaving of controller ports for multipath configuration is not supported.
- To connect a single enclosure, cable the ports so that the enclosure is connected to port 0 and 2, or port 1 and 3 on the controller card.
- To connect multiple enclosures, connect a set of ports on the controller to the first enclosure as mentioned in the above bullet point. Connect cables from the first enclosure to the second enclosure. Then, continue to connect 2 cables from the last enclosure in the chain to the next enclosure you want to include until you have up to 4 enclosures, or 8 enclosures with 4 enclosures per stack. See Figure 6 and Figure 7.

NOTE: The PERC H965e adapter card can detect and use redundant paths to disks contained in the enclosures. With redundant paths to the same device, if one path fails, the other path can be used to communicate between the controller and the device.

- NOTE:** Max enclosure support for multipath: Eight (8) MD2412 and MD2424 enclosures or four (4) MD2460 enclosures.
- NOTE:** It is recommended that the wait-time for hot-removing or hot-inserting enclosure data or power cables is 2-3 minutes for 24 drives in a single MD24xx series disk-attached storage enclosure and up to five minutes for maximum enclosures configuration.




Figure 7. PERC H965e adapter card ports—Eight enclosures with multipath configuration for the MD2412 and MD2424 series enclosures




Figure 8. PERC H965e adapter card ports—Four enclosures with multipath configuration for the MD2460 series enclosures

Virtual disk features

NOTE: All VD's are exposed as 4k sector devices to the host and operating system.

This section lists the following virtual disk features supported on PERC 12 cards:

- [Auto-import foreign virtual drives](#)
- [Virtual disk migration](#)
- [Virtual disk write cache and read-ahead policy](#)
- [Virtual disk initialization](#)
- [Online capacity expansion](#)
- [Background operations](#)

Auto-import foreign virtual drives

By default, this feature is enabled on a PERC 12 controller. Foreign Virtual Drives (VDs) are automatically imported when a server is either restarted, or a controller is reset or replaced. Ensure that you verify the new VD number after automatically importing because a VD number may change after an auto-import operation. To resolve any issues pertaining to auto-importing, see the Troubleshooting section in this User's Guide.

Virtual disk migration

The PERC 12 series supports migration of virtual disks from one controller to another without taking the target controller offline. The controller can import RAID virtual disks in optimal, degraded, or partially degraded states. You cannot import a virtual disk that is offline. When a controller detects a configured physical disk, it marks the physical disk as foreign, and generates an alert indicating that a foreign disk was detected.

NOTE: You cannot migrate data managed by earlier versions such as PERC 10 and PERC 11 to PERC 12 controllers. However, you can migrate data of unsecured non-RAID disks from PERC 10 and PERC 11 to the latest PERC 12 controllers.

NOTE: The source controller must be offline prior to performing the disk migration.

NOTE: Importing secured VD's is supported as long as the appropriate Local Key Management (LKM) is supplied or configured.

NOTE: An additional reboot operation may be required for the UEFI BIOS to detect a newly imported Virtual disk.

NOTE: Importing a manually offlined VD is not supported. Recreate the VD exactly as it was before it was offline to get the configuration back.

Online capacity expansion

You can dynamically expand an online VD by using Online Capacity Expansion (OCE).

NOTE: The OCE feature is not supported on spanned VD's such as RAID 50 and RAID 60. Also, OCE is not supported on RAID 10.

NOTE: If multiple VD's exist within a common disk group or if a single VD does not start at the first block of the disk group, the OCE operation is not supported on those VD's.

NOTE: OCE typically impacts disk performance until the expansion operation is completed.

Drive mixing rules apply to an OCE operation which does not allow mixing of:

- SAS and SATA drives
- HDDs and SSDs
- NVMe drives with HDDs and SSDs

An OCE operation can be performed in the following methods:

- If there is free space available in a disk group, the capacity of a virtual disk can be expanded within the available free space.

- If there is no free space available within the disk group, physical disks can be added to the virtual disk to increase its capacity.
- After replacing all disk group members with larger drives than the original member drives, use the PERC CLI utility to expand the existing disk group to a larger size using the Expand Disk Group feature.

⚠ CAUTION: Do not remove drives or attempt disk migration during OCE operations. This results in loss of VDs.

i NOTE: An OCE operation is not supported when another OCE operation with PD addition is in progress on the controller.

i NOTE: When an OCE operation is started on a VD, all VDs with write-back cache policy that is present on the controller are converted to write-through. Cache policy is restored to its original setting after the OCE operation is completed.

i NOTE: An OCE operation is not supported when certain background operations are in progress on a VD or conversely. Any pending operations will start after the current operations are completed. Example operations are:

- Virtual Disk Initialization
- Consistency Check
- Background Initialization
- Drive Rebuild
- Copyback
- Patrol Read

i NOTE: If the controller already contains the maximum number (64) of VDs (all VD configurations except a single drive R0), contains the maximum number (240) of disk groups, or surpasses the maximum number (240) of configured PDs with the OCE operation, you cannot perform an OCE operation.

i NOTE: The RAID Level Migration (RLM) feature is not supported on PERC12 controllers.

Virtual disk write cache and read-ahead policy

The write cache policy of a virtual disk determines how the controller handles Write operations to the virtual disk.

Table 5. Write cache policies

Feature	Description
Write-back	The controller sends a data transfer completion response to the host when the controller cache has received all the data in a transaction. The controller then writes the cached data to the storage device in the background. i NOTE: The default cache setting of HDD-based VDs is write-back. Write-back caching is also supported for single drive RAID 0 virtual disks.
Write-through	The controller sends a data transfer completion response to the host system when the disk subsystem has received all the data in a transaction. i NOTE: Certain data patterns and configurations perform better with a write-through cache policy.
Read-Ahead	Available in 8.8.0.0.18-26 and later versions of PERC12 controllers. The controller speeds up sequential reads on virtual disks. This is supported only on HDDs with write-back cache policy. i NOTE: You can request for Read Ahead support on any VD including SSD but controller will change only the cache setting if it is supported.

i NOTE: All RAID volumes are presented as write-through to the operating system (Windows and Linux) independent of the actual write cache policy of the virtual disk. PERC cards manage the data in cache independently of the operating system or any applications.

i NOTE: Use the Dell OpenManage Storage Management application or the HII Configuration Utility to view and manage virtual disk cache settings.

Conditions under which write-back is used

Write-back caching is used under all conditions in which the battery (Energy Pack) is present and in good condition.

Conditions under which forced write-back with no battery is used

 **CAUTION:** It is recommended that you use a power backup system when forcing write-back to ensure there is no loss of data if the system suddenly loses power.

The write-back mode is available when you select force write-back with no battery (Energy Pack). When the forced write-back mode is selected, the virtual disk is in write-back mode even if the battery is not present.

Virtual disk initialization

The PERC 12 series controllers support two types of virtual disk initialization:

- Full initialization
- Fast initialization

 **CAUTION:** Initializing virtual disks erases files and file systems while keeping the virtual disk configuration intact.

Full initialization

 **NOTE:** When the full initialization operation is in progress, you cannot perform any host data I/O operations on the VD.

Performing a full initialization operation on a virtual disk overwrites all blocks and deletes any data that previously existed on the virtual disk. Full initialization of a virtual disk eliminates the need for the virtual disk to undergo a Background Initialization (BGI). Full initialization can be performed after the virtual disk is created. For more information on how to create a virtual disk and configure virtual disk parameters (full initialization) in HII, see [Create virtual disks](#).

 **NOTE:** If the system reboots during a full initialization operation, the operation is abruptly stopped and a BGI operation starts on the virtual disk.

Fast initialization

A fast initialization on a virtual disk overwrites the first and last 8 MiB of the virtual disk, clearing any boot records or partition information. The operation takes only 2–3 seconds to complete, but it is followed by BGI, which takes longer to complete.

 **NOTE:** During full or fast initialization, the host cannot access the virtual disk. As a result, if the host attempts to access the virtual disk while it is initializing, all I/O sent by the host will fail.

 **NOTE:** When using iDRAC to create a virtual disk, the drive undergoes fast initialization. During this process all I/O requests to the drive will respond with a sense key of **Not Ready** and the I/O operation will fail. If the operating system attempts to read from the drive as soon as it discovers the drive and while the fast initialization is still in process, the I/O operation fails and the operating system reports an I/O error.

Background operations

Background initialization

Background Initialization (BGI) is an automated process that writes parity or mirror data on newly created virtual disks. BGI does not run on RAID 0 virtual disks. You can control the BGI rate in the HII application. Any change to the BGI rate does not take effect until the next BGI operation is started.

 **NOTE:**

- If you cancel BGI, it automatically restarts within five minutes.

- Unlike full or fast initialization of virtual disks, background initialization does not clear data from the physical disks.
- Consistency Check (CC) and BGI typically cause some loss in performance until the operation completes.

Consistency check and BGI perform similar functions in that they both correct parity errors. However, CC reports data inconsistencies through an event notification, while BGI does not. You can start CC manually, but not BGI.

Consistency checks

Consistency Check (CC) is a background operation that verifies and corrects the mirror or parity data for fault tolerant virtual disks. It is recommended that you periodically run a consistency check on virtual disks. You can manually start a CC using the HII Configuration Utility or the PERC CLI application. To start a CC using the HII Configuration Utility, see [Check consistency of VDs](#).

NOTE: CC or BGI typically causes some loss in performance until the operation completes.

CC and BGI both correct parity errors. However, CC reports data inconsistencies through an event notification, while BGI does not. You can start CC manually, but not BGI.

Drive features

This section lists the following hard drive features supported on PERC 12 cards:

- [Self-Encrypting Drives \(SEDs\)](#)
- [Opal Security Management](#)
- [Instant secure erase](#)
- [4KB sector drives](#)
- [Non-Volatile Memory Express](#)
- [Conditions under which a PERC supports an NVMe drive](#)

Self-Encrypting Drives (SEDs)

The PERC 12 series of cards support SEDs for protection of data against loss or theft by the use of encryption technology on the drives. There is one security key per controller. You can manage the security key using Local Key Management (LKM) or OpenManage Secure Enterprise Key Manager also referred as Secure Enterprise Key Manager (SEKM). The security key is used by the controller to lock and unlock access to encryption-capable physical disks. To use this feature, you must:

- Have SEDs in your system.
- Create a security key.

SEDs that are secured by a non-PERC entity cannot be used by PERC. Ensure that the SED is re-provisioned in an applicable manner by the non-PERC entity before connecting to PERC. For more information, see [Security key and RAID management](#).

NOTE: You cannot enable security on non-optimal virtual disks.

NOTE: PERC 12 supports Trusted Computing Group (TCG) Enterprise Security Subsystem Classes (SSC) SAS or SATA SEDs, and TCG Opal SSC NVMe drives.

Enhanced security for SEDs

In Self-Encrypting Drives (SEDs) managed by controllers 8.8.0.0.18-26 and later versions, all existing and newly created SEDs are secured by the following features. When such drives are erased, setting on these drives are reverted to their default.

- Disables ports used for downloading firmware and diagnostic analysis—Some SEDs have special access ports that will allow or block operations if the drive is locked. If a secured drive has these ports, the firmware will disable on such drives. When the ports are disabled, firmware update and retrieval of diagnostic information is blocked if a drive is in the locked state.
- Disables SED manufacturers from accessing data for issue analysis— Some SEDs enable its manufacturers to access data for troubleshooting purposes. However, this will no longer be allowed and manufacturers access is disabled.
- Full administrator authority ownership— All SEDs will have their admin credentials changed to the security key. Any drive with the admin credentials changed will not be able to be erased using normal cryptographic erase if the drive is in a locked state. Any drive in locked state will require a PSID erase using the PSID key listed on the drive label. Drives cannot be

migrated to older versions of firmware as the older firmware does not know how to manage the changed Admin credentials. Firmware downgrade to previous versions while secured drives are present will be blocked.

Opal Security Management

Opal Security Management of Opal SEDs requires security key management support. You can use integrated Dell Remote Access Controller (iDRAC) to generate the security key for the Opal drives which is used as an authentication key to lock and unlock the Opal drives.

Instant secure erase

Instant Secure Erase (ISE) drives use the same encryption technology as SEDs but do not allow the encryption key to be secured. The encryption technology allows the drive to be re-purposed and securely erased using the cryptographic erase function.

i **NOTE:** ISE drives do not provide protection against reading the drive data in case of a theft.

Cryptographic Erase with PSID Revert

The controller firmware supports reverting an SED state to the factory default with the help of Physical Security Identifier (PSID). The PSID, or Physical SID of the drive, is a 32-character password assigned by the drive manufacturer during production. A host system cannot change the password. The PSID is on the drive label in a readable format, and depending on the drive manufacturer, it may also be available in a bar code format. If a drive is in the locked state and the key is not known then a cryptographic erase with PSID is required. The SED physical security ID (PSID) is required for reverting an SED to the non-secured factory state. For enhanced security, the PSID is accessible only by removing the drive and examining the drive label. In addition to reverting the drive state to factory default settings, it also secure erases all existing data.

i **NOTE:** The PSID can only be used for reverting the drive; it does not grant access to any encrypted data present on the drive.

Physical disk erase

The Physical disk erase feature allows data that is saved on disks to be securely erased so that data cannot be recovered. The H965 series cards provide the following four methods for erasing data on drives provided the drive supports it.

- Physical disk erase—Erases drives by writing data pattern on disks with varying number of passes. Physical disk erase is supported on non-ISE and non-SEDs.
- Cryptographic Erase—Cryptographically erases disks by changing the media encryption key. This feature is supported on ISE or SEDs. In case of SEDs, unsecures the drives and reverts them to factory security settings. See [Instant Secure Erase](#) and [Self-Encrypting Drives \(SEDs\)](#).
- Sanitize Block Erase—Alters information by setting the physical blocks to a vendor specific value.
- Sanitize Overwrite Erase—Alters information by setting physical blocks to a user-specific value. Also, it can do multiple overwrite operations and invert the pattern between consecutive overwrite passes.

Sanitize Erase operation

The Sanitize Block Erase or Sanitize Overwrite Erase operation managed by the controllers can be triggered using PERC CLI. Prerequisites for performing this operation are:

- Have an unconfigured physical disk.
- Physical disks must support only the necessary sanitize method.

i **NOTE:** If drive supports the cryptographic erase operation, drive will not be eligible to perform the required sanitize operation. Limitations of the PERC managed Sanitize Erase operation are:

- After the sanitize erase operation is started, the non-RAID drive creation operation will be blocked while an erase operation is in progress.

4KB sector drives

The PERC 12 series controllers support 4 KB sector drives. Before installing Windows on 4KB sector drives, see [Windows operating system installation errors](#).

SSD based VD's are presented to the hypervisor as a hard drive. VD's that are a part of SSDs or presented to ESXi 8.x and earlier versions continue to be enumerated as hard drive in ESXi 9.0 and later versions.

NOTE: Mixing 512-byte native and 512-byte emulated drives in a VD is allowed, but mixing 512-byte and 4 KB native drives in a VD is not allowed.

NOTE: All VD's are exposed to the host as a 4K device.

Non-Volatile Memory Express

Non-Volatile Memory Express (NVMe) is a standardized, high-performance host controller interface and a storage protocol for communicating with non-volatile memory storage devices over the PCIe interface standard. The PERC 12 controller supports up to 8 direct-attach NVMe drives. The PERC 12 controller is a PCIe endpoint to the host, a PowerEdge server, and configured as a PCIe root complex for downstream PCIe NVMe devices connected to the controller.

NOTE: The NVMe drive on the PERC 12 controller is identified as a SCSI disk in the operating system. The NVMe CLI will not work for the attached NVMe drives.

Conditions under which a PERC supports an NVMe drive

- A single Namespace must be present.
- The NameSpace Identifier (NSID) with ID 1, which is (NSID=1), must be present.
- The namespace with NSID=1 must be formatted without end-to-end data protection information and must have the metadata disabled
- PERC supports 512-bytes or 4 KB sector disk drives for NVMe devices.

Drive recovery for NVMe initialization failure

If an NVMe drive fails to initialize, the drive that is connected to PERC can be recovered in HII. The NVMe initialization errors in the drives are listed as recoverable and nonrecoverable errors in HII.

The recoverable drives are identified as the following in the HII: `Unusable (Recoverable), Error: <error string>`. The nonrecoverable drives are listed as `Usable` in HII. Metadata and End-to-end Data Protected (EEDP) formatted disks cannot be recovered when connected to PERC 12.

Drive recovery with correctable NVMe initialization errors

Recover the drives with recoverable NVMe initialization errors in HII to enable the drives to work properly.

Steps

1. Log in to HII.
2. Click **Main Menu > Device Management > Enclosure X**.
The drives with recoverable and non-recoverable errors are listed.
3. Select the drive and click **Recover**.
If the drive data is successfully recovered, the drive is listed under physical drives and removed from the recoverable error list. If the drive has other recoverable errors, the drive is listed again in the recoverable errors list.
4. If the repair is not successful, click **Recover**.
If the error is still not resolved or if the drive has other non-recoverable errors, the drive is moved to the non-recoverable error list.

Fault tolerance

The Controller series cards support the following:

- [The SMART feature](#)
- [Patrol Read](#)

The following sections describe methods to achieve fault tolerance.

The SMART feature

The SMART feature monitors certain physical aspects of all motors, heads, and physical disk electronics to help detect predictable physical disk failures. Data on SMART compliant physical disks can be monitored to identify changes in values and determine whether the values are within threshold limits. Many mechanical and electrical failures display some degradation in performance before failure.

A SMART failure is also referred to as predicted failure. There are numerous factors of a predicted physical disk failure—bearing failure, a broken read or write head, and changes in spin-up rate. Also, there are factors related to read or write surface failure, such as seek error rate and excessive bad sectors.

 **NOTE:** For detailed information on SCSI interface specifications, see t10.org. For detailed information on SATA interface specifications, see t13.org.

Patrol Read

The Patrol read feature is designed as a preventative measure to ensure physical disk health and data integrity. Patrol read scans and resolves potential problems on configured physical disks. The Dell OpenManage Storage Management application can be used to start patrol-read and change its behavior.

Some properties of the patrol-read feature are:

- Runs on all disks on the controller that are configured as part of a virtual disk, including hot-spares.
- The Patrol read feature does not run on physical disks that are not part of non-RAID disks, virtual disks, or drives that are in Ready state.
- The amount of controller resources dedicated to patrol read operations adjusts based on the number of outstanding disk I/O operations. For example, if the system is processing a large number of I/O operations, then patrol read uses fewer resources to allow the I/O to take a higher priority.
- Patrol read does not run on disks that are involved in any of the following operations:
 - Rebuild
 - Replace member
 - Full or background initialization
 - Consistency Check
 - Online capacity expansion (OCE)

 **NOTE:** By default, patrol read automatically runs every seven days on configured hard drives.

For more information about patrol read, see the *Dell OpenManage Storage Management User's Guide* available at [OpenManage Manuals](#).

Physical disk failure detection

If a disk fails and it is replaced with a new disk, the controller will automatically start a rebuild on the new disk. See, [Configured slot behavior](#). Automatic rebuilds can also be performed on hot-spares. If you have configured hot-spares, the controller will automatically try to use them to rebuild the degraded virtual disk.

Using persistent hot-spare slots

The PERC 12 series is configured so that the system backplane or storage enclosure disk slots are dedicated as hot-spare slots.

Any slots with hot-spares are persistent. If a hot-spare disk fails or is removed, a replacement disk that is inserted into the same slot automatically becomes a hot-spare with the same properties as the hot-spare slot it is replacing. If the replacement disk does not match the disk protocol and technology, it does not become a hot-spare.

For more information about persistent hot-spares, see the Dell OpenManage documentation available on the support site.

- NOTE:** If all the VDs that are associated with a dedicated hot-spare are removed, then the hot spare will be converted to a global hot-spare.
- NOTE:** If all the VDs attached to a controller are removed from a system, then the hot-spare slots, if any, will be deleted.

Configured slot behavior

This feature is similar to persistent hot-spare slot behavior. If a fault-tolerant VD is configured to the system and a drive is replaced, the configured slot will automatically rebuild (copyback) on the inserted drive regardless of the data on the drive. This operation overwrites the data on the drive.

Table 6. Drive state operation

Drive state/operation	Drive state	Slot configured in VD
Insert unconfigured drive into the system.	Ready	Rebuild or copyback starts
Insert configured drive into the system.	Foreign	<ul style="list-style-type: none"> ● Rebuild or copyback starts ● Original drive data lost
Insert configured locked drive into the system (unlockable).	Foreign	Cryptographic Erase (If configured VD is not secured). <ul style="list-style-type: none"> ● Rebuild or copyback starts ● Original drive data lost
Insert locked drive into the system (non-unlockable).	Foreign locked	Foreign locked

Physical disk hot-swapping

Hot swapping is the manual replacement of a disk while the PERC 12 series cards are online and performing their normal functions. Ensure that the following requirements are met before hot-swapping a physical disk:

- The system backplane or enclosure must support hot-swapping.
 - To perform a rebuild or copyback operation, the replacement drive must be of the same protocol and disk technology. For example, only a SAS hard drive can replace a SAS hard drive and only a SATA SSD can replace a SATA SSD.
- NOTE:** To check if the backplane supports hot-swapping, see the platform-specific Installation and Service Manual available on the support site.

Using Replace Member and Revertible hot-spares

The Replace Member functionality enables a previously commissioned hot-spare to revert to a usable hot-spare. When a disk failure occurs within a virtual disk, an assigned hot-spare—dedicated or global—is commissioned and begins rebuilding until the virtual disk is optimal. After the failed disk is replaced in the same slot and the rebuild to the hot-spare is complete, the controller automatically starts to copy data from the commissioned hot spare to the newly inserted disk. After the data is copied, the new disk is a part of the virtual disk and the hot-spare is reverted to being a ready hot-spare. This allows hot-spares to remain in specific enclosure slots. While the controller is reverting the hot-spare, the virtual disk remains optimal. The controller automatically reverts a hot-spare only if the failed disk is replaced with a new disk in the same slot. If the new disk is not placed in the same slot, a Manual Replace Member operation can be used to revert a previously commissioned hot-spare.

- NOTE:** A Manual Replace Member operation typically causes a temporary impact on disk performance. After the operation is complete, performance returns to normal state.
- NOTE:** Replace member is also referred to as "copyback and replace" in some management applications or events.

Automatic Replace Member with predicted failure

An Automatic Replace Member feature is triggered when there is a SMART predictive failure reported on the VD of a PD. The Automatic Replace Member feature is initiated when the first SMART error occurs on a physical disk that is part of a virtual disk. The target disk needs to be a hot-spare that qualifies as a rebuild disk. The physical disk with the SMART error is marked as failed only after the successful completion of the replacement task. This prevents the array from reaching degraded state.

If an Automatic Replace Member operation occurs using a source disk that was originally a hot-spare (that was used in a rebuild), and a new disk is added and set as a target disk for the replace member operation, the hot-spare drive will revert to the hot-spare state after the operation completes.

 **NOTE:** To enable the Automatic Replace Member feature, use the Dell OpenManage Storage Management application.

Controller cache

The PERC 12 series of cards contain local DRAM on the controllers. This DRAM can cache I/O operations for improvement of both read and write operations.

I/O workload to HDDs, such as random read-write 512 B and 4 kB, may take some time to flush cached data. Cache is flushed periodically, but for configuration changes or system shutdown, the cache is required to be flushed before the operation can be completed. It can take several minutes to flush cache for some workloads depending on the speed of the HDDs and the size of data in the cache.

The following operations require a complete cache flush:

- Configuration changes (add or delete VDs, VD cache setting changes, foreign configuration scan, and import)
- System restart or shutdown
- Abrupt power loss causing [Controller cache preservation](#).

 **NOTE:** iDRAC or OpenManage Storage Management periodically scans for the foreign configurations when the foreign disks are present. This action degrades the performance. If a foreign disk is present, it is recommended that you import, clear, or remove the foreign disk to prevent an impact on the performance.

Controller cache preservation

The controller is capable of preserving its cache even if the server is improperly shut down or if a power outage occurs. The PERC 12 series controller is attached to a Battery Backup Unit (BBU) that provides backup power during system power loss to preserve the controller's cache data.

Cache preservation with non-volatile cache

The Non-Volatile Cache (NVC) allows controller cache data to be stored indefinitely. If the controller has data in the cache memory during a power outage or improper system shutdown, little power from the battery (Energy Pack) is used to transfer the cache data to non-volatile flash storage where it remains until power is restored and the system is started. If the cache preservation process is interrupted by power-on, the controller may request an extra reset during the boot to complete the process.

Recovering cache data

About this task

To recover cached data in case of power loss or improper shutdown of the server, do the following:

Steps

1. Restore the system power.
2. Start the system.
3. If preserved cache exists on the controller, a message is displayed. For more information about how to recover cache, see [Preserved Cache state](#).

Battery (Energy Pack) Transparent Learn Cycle

A Transparent Learn Cycle (TLC) is a periodic operation that calculates the charge that is remaining in the battery (Energy Pack) to ensure that there is sufficient energy. The operation runs automatically and causes no impact to the system or controller performance.

The controller automatically performs a TLC operation on the battery (Energy Pack) to calibrate and gauge its charge capacity every 90 days. The operation can be performed manually, if required.

NOTE: If enabled, virtual disks stay in the write-back mode during a TLC. When the TLC operation completes, the controller sets the next TLC to 90 days in the future.

NOTE: During a Transparent Learn Cycle, the battery will be discharged to around 70% to calibrate and gauge the capacity. After the cycle is completed, the battery will charge back up to full capacity.

Transparent Learn Cycle completion time

The time frame for completion of a learn cycle is a function of the battery (Energy Pack) charge capacity and the discharge and charge currents used. Typical time completion for a transparent learn cycle is between 4 to 8 hours. If the learn cycle is interrupted, it begins at a new cycle.

Conditions for replacing the battery (Energy Pack)

A PERC card battery (Energy Pack) is indicated as Failed when the battery stops functioning. If the battery is declared as failed, then all the virtual disks in write-back mode transitions to write-through mode. After replacing the battery, virtual disk transitions back to the write-back mode.

Install and remove a PERC 12 card

Topics:

- [Before working inside your system](#)
- [Safety instructions](#)
- [Remove the PERC H965i Adapter](#)
- [Install the PERC H965i Adapter](#)
- [Remove the PERC H965i Front card](#)
- [Install the PERC H965i Front card](#)
- [Remove PERC H965i MX](#)
- [Install PERC H965i MX](#)
- [Removing the PERC H965e Adapter card](#)
- [Installing the PERC H965e Adapter card](#)

Before working inside your system

Prerequisites

Follow the safety guidelines listed in [Safety instructions](#)

Steps

1. Power off the system and all attached peripherals.
2. Disconnect the system from the electrical outlet and disconnect the peripherals.
3. If applicable, remove the system from the rack.
For more information, see the Rail Installation Guide relevant to your rail solutions at [PowerEdge Manuals](#).
4. Remove the system cover.

Safety instructions

 **CAUTION:** Ensure that two or more people lift the system horizontally from the box and place it on a flat surface, rack lift, or into the rails.

 **WARNING:** Opening or removing the PowerEdge server cover while the server is powered on may expose you to a risk of electric shock.

 **WARNING:** Do not operate the server without the cover for a duration exceeding five minutes. Operating the system without the system cover can result in component damage.

 **NOTE:** Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

 **CAUTION:** To ensure proper operation and cooling, all system bays and fans must always be populated with a component or a blank.

 **NOTE:** It is recommended that you always use an antistatic mat and antistatic strap while working on components inside the server.

NOTE: To ensure proper operation and cooling, all system bays and fans must always be populated with a component or a blank.

NOTE: While replacing the hot swappable PSU, after next server boot, the new PSU automatically updates to the same firmware and configuration of the replaced one.

Remove the PERC H965i Adapter

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

NOTE: Tasks indicated in this procedure may vary based on the server model from which the PERC H965i Adapter is removed.

Steps

1. Power off the server, including any attached peripherals, and disconnect the system from the electrical outlet and peripherals.

NOTE: Perform a graceful shutdown of the system to ensure data in the cache is moved to the disk(s) before the controller is removed.

2. Open the system.
3. Locate the PERC card in the expansion riser on the system board.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

4. Unfasten and lift the riser from the system board. Remove the PERC card.
5. Unfasten the screw that is used to hold the cable cover.
6. Disconnect any cables connected to the card:
 - a. Press down and hold the metal tab on the cable connector.
 - b. Pull the cable out of the connector.
7. Replace the storage controller card and reconnect the data cables before placing them in the riser. Tighten the screw again to hold the cable cover firmly. For more information on installing the card, see [Install the PERC H965i Adapter](#).
8. Reinstall the riser on the system board and fasten the riser.
9. Close the system.
10. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.




Figure 9. Remove the PERC H965i Adapter

Install the PERC H965i Adapter

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

NOTE: The steps may vary based on the platforms where PERC H965i Adapter is installed.

Steps

1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.
2. Open the system.
3. Align the card-edge connector with the connector on the riser.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

4. Press the card-edge until the card is fully seated.
5. Unfasten the screw that is used to hold the cable cover.
6. Connect the data cable connectors to the card. Tighten the screw again to hold the cable cover firmly.
7. Route the data cable through the channel on the inner side of the server chassis to the backplane.
8. Attach the connector to the corresponding connector on the backplane as labeled on the controller. If you are installing a PERC 12 H965i adapter card that has a detachable battery, for example, in a PowerEdge C6620 sled, install the battery in the air shroud by doing the following:
 - a. Insert the battery into the air shroud by ensuring that the surface that has the battery part number label is on top.

- b. Carefully place the battery cable in the air shroud notch and push the battery down. Hooks on both sides of the battery lock it into the air shroud slots and the battery is firmly held in position.
 - c. Connect the battery socket to the controller socket.
 - d. Install the air shroud. See the "Installing the air shroud" section in the *Installation and Service Manual* of the respective sled model available on the support site.
 - e. Install the riser that holds the PERC 12 card. See the "Installing the expansion card riser" section in the *Installation and Service Manual* of the respective sled model available on the support site.
 - f. Carefully route the battery cable underneath the heat sink in the space available between the DIMMs and CPU socket.
9. Close the system.
 10. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.




Figure 10. Install the PERC H965i Adapter

Remove the PERC H965i Front card

The PERC H965i Front controller card is available in two types. 1) The PERC H965i Front controller card is supported on 16th Generation (16G) PowerEdge servers. 2) The PERC H965i Front DC-MHS controller card is supported on 17th (17G) Generation PowerEdge servers. Both controller types have separate Update Packages (DUPs). However, the procedure for installation and removal of both the controller cards is same.

Prerequisites

The PERC H965i Front controller card is available in two types:

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and anti-static wrist strap while working on components inside the system.

Steps

1. Power off the server, including any attached peripherals, and disconnect the system from the electrical outlet and peripherals.

NOTE: Perform a graceful shutdown of the system to ensure data in the cache is moved to the disk before the controller is removed.

2. Open the system.
3. Locate the PERC card in the controller carrier at the front of the system.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

4. Unscrew the fasteners on the controller carrier and slide the carrier away from the backplane, disconnecting the controller from the backplane. When you unfasten the four (4) controller screws, even the plastic shroud that holds the battery is removed.
5. Disconnect any cables connected to the card:
 - a. Press down and hold the metal tab on the cable connector.
 - b. Pull the cables out of the connector.
6. Remove the PERC controller from the controller carrier.
 - If necessary, remove the battery from the plastic shroud by carefully pulling it out, and then removing the cable that connects the battery to the controller.
 - To replace the battery, connect the battery cable to the controller, place the battery between the guiding ribs on the plastic shroud.
 - Slowly press-fit to lock it firmly.
7. Insert the replacement controller into the carrier and secure it with the appropriate screws.
8. Take the replacement storage controller and reconnect the cables before reconnecting it to the backplane.
9. Close the system.
10. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.




Figure 11. Remove the PERC H965i Front card

Install the PERC H965i Front card

The PERC H965i Front controller card is available in two types. 1) The PERC H965i Front controller card is supported on 16th Generation (16G) PowerEdge servers. 2) The PERC H965i Front DC-MHS controller card is supported on 17th (17G) Generation

PowerEdge servers. Both controller types have separate Update Packages (DUPs). However, the procedure for installation and removal of both the controller cards is the same.

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and anti-static wrist strap while working on components inside the system.

Steps

1. Power off the system, including any attached peripherals, and disconnect the system from the electrical outlet.

NOTE: Perform a graceful shutdown of the sled to ensure that data in the cache is moved to the disk before the controller is removed.

2. Open the system.

3. Connect the PERC card to the carrier and ensure that the screws are properly fastened in place. For information about removing or replacing the battery, see [Remove the PERC H965i Front card](#).

CAUTION: To prevent damage to the card, hold the card by its edges only.

4. Connect the cable connectors to the card.

NOTE: Ensure that you connect the cable according to the connector labels on the cable. The cable does not function properly if reversed.

5. Align the carrier with the guide pins until the controller is securely seated.

6. Slide the card into the connector until it is fully seated in the connector. Tighten the screws on the carrier that connect to the chassis to secure the carrier.

7. Close the system.

8. Reconnect the system to its electrical outlet and power on the system and any attached peripherals.




Figure 12. Install the PERC H965i Front card

Remove PERC H965i MX

Prerequisites

 **CAUTION:** Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

 **CAUTION:** To prevent damage to the card, hold the card by its edges only.

 **NOTE:** It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

 **NOTE:** The steps may vary based on the platform from where PERC H965i MX is removed.

Steps

1. Turn off the sled, including any attached peripherals, and remove the sled from the MX chassis.

 **NOTE:** Perform a graceful shutdown of the system to ensure that data in the cache is moved to the disk before the controller is removed.

2. Open the sled.

3. Locate the PERC card on the system board.

 **CAUTION:** To prevent damage to the card, hold the card by its edges only.

4. Using the blue tab, rotate the lever of the controller.

5. Pull the release lever upward to disengage the controller from the connector.

6. Disconnect the cable from the card. To disconnect the cable:

- a. Press and hold the metal tab on the cable connector.
- b. Pull the cable out of the connector.

7. Lift the card from the system board.

8. Replace the storage controller card and connect the cable. For information on installing the card, see [Install PERC H965i MX](#).

9. Close the sled.

10. Insert the sled into the MX chassis and turn on the system and any attached MX chassis peripherals.




Figure 13. Remove the PERC H965i MX

Install PERC H965i MX

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

NOTE: Steps may vary based on the platform where PERC H965i MX is installed.

Steps

1. Turn off the sled and any attached peripherals, and remove the sled from the MX chassis.
2. Open the sled.
3. Connect the backplane data cable connector to the card.

NOTE: Ensure that you connect the cable according to the connector labels on the cable. The cable does not function properly if reversed.

4. Align the bracket notches with the tabs on the sides of the sled chassis and align the PERC card connector with the connector on the system board.

CAUTION: To prevent damage to the card, hold the card by its edges only.

5. Press the PERC card into the connector until it is firmly seated.
6. Press the release lever to secure the card to the sled.

NOTE: The pin on the release lever secures the card to the chassis of the sled.

7. Route the data cable through the clip on the card and through the channel on the inner side of the chassis.
8. Attach the connector to the corresponding connector on the backplane as labeled in the controller.
9. Close the sled.
10. Insert the sled into the MX chassis and turn on the system and any attached MX chassis peripherals.




Figure 14. Install the PERC H965i MX

Removing the PERC H965e Adapter card

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

1. Power off the system, including any attached peripherals, and disconnect the system from the electrical outlet and peripherals.
2. Open the system.
3. Locate the PERC card on the system board.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

4. Disconnect the SAS4 SFF cables, if connected, to any external storage enclosure(s)
5. Remove the corresponding riser and lift the card from the PCIe slot.

6. Replace the storage controller card and connect the cable.
7. Close the system.
8. Reconnect the system to its electrical outlet and power on the system, including any attached peripherals.




Figure 15. Remove the H965e Adapter card

Installing the PERC H965e Adapter card

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

1. Power off the system, including any attached peripherals, and disconnect the system from the electrical outlet.
2. Open the system.
3. Align the card-edge connector with the PCIe slot on the system riser.

CAUTION: To prevent damage to the card, hold the card by its edges only.

4. Press the card-edge down until the card is fully seated in the PCIe slot on the system riser.
5. Insert the riser into the system board and connect the SAS4 SFF cables to the card.
6. Close the system.
7. Reconnect the system to its electrical outlet and power on the system and any attached peripherals.




Figure 16. Install the PERC H965e Adapter card

Driver support for PERC H965i Adapter, PERC H965i Front, PERC H965i MX, and PERC H965e Adapter

The PERC H965i Adapter, PERC H965i Front, PERC H965i MX, and PERC H965e Adapter cards require software drivers to operate with the supported operating systems.

This chapter contains the procedures for installing the drivers for the H965 series controller cards. The two methods for installing a driver that is discussed in this chapter are:

- **Installing a driver during operating system installation:** Use this method if you are performing a new installation of the operating system and want to include the drivers.
- **Updating existing drivers:** Use this method if the operating system and the RAID controllers are already installed and you want to update to the latest drivers.

 **NOTE:** It is recommended to download the latest Out-of-Box (OoB) drivers from the Dell Support site because the Inbox drivers in the operating system may not contain the full functionality and the latest fixes.

Topics:

- [Creating the device driver media](#)
- [Windows driver installation](#)
- [Update PERC H965i Adapter, PERC H965i Front, PERC H965i MX, and PERC H965e Adapter driver that runs on Microsoft Windows operating system](#)
- [Linux driver installation](#)
- [Load the driver while installing an operating system](#)

Creating the device driver media

Use one of the following two methods to create the device driver media:

- [Download and save PERC H965i Adapter, PERC H965i Front, PERC H965i MX, and PERC H965e Adapter, drivers from the support site](#)
- [Download and save PERC H965i Adapter, PERC H965i Front, PERC H965i MX, and PERC H965e Adapter drivers from the Dell Systems Service and Diagnostic Tools](#)

Download and save PERC H965i Adapter, PERC H965i Front, PERC H965i MX, and PERC H965e Adapter, drivers from the support site

Download the latest Out-Of-Box (OOB) drivers from the Dell Support site because the Inbox drivers in the operating system may not contain the full functionality and the latest fixes.

About this task

To download drivers from the Dell Support website, do the following:

Steps

1. Go to the support site.
2. Enter the Service Tag of your system in the **Choose by Service Tag to get started** box or select **Choose from a list of all Dell products**.
3. Select the system type, operating system, and the SAS RAID category from the drop-down menu.

The drivers that are applicable to your selection are displayed.

4. Download the required drivers to a USB drive, CD, or DVD.
5. When installing an OS, use the media that you created to load the driver. For more information on reinstalling the operating system, see the operating system documentation at [Operating System Documentation](#).

Download and save PERC H965i Adapter, PERC H965i Front, PERC H965i MX, and PERC H965e Adapter drivers from the Dell Systems Service and Diagnostic Tools

About this task

To download drivers from the **Dell Systems Service and Diagnostic Tools** media:

Steps

1. Insert the **Dell Systems Service and Diagnostics Tools** media in your system. The **Welcome to Dell Service and Diagnostic Utilities** screen is displayed.
2. Select your system model and operating system.
3. Click **Continue**.
4. From the list of drivers displayed, select the driver that you require.
5. Select the self-extracting ZIP file and click **Run**.
6. Copy the driver to a CD, DVD, or USB drive.
7. Repeat steps 1 to 6 for all the drivers that you require.

Windows driver installation

Before you install the Windows driver for PERC H965i Adapter, PERC H965i Front, PERC H965i MX, and PERC H965e Adapter, you must first create a device driver media.

- Read the Microsoft **Getting Started** document that shipped with your operating system.
- Ensure that your system has the latest BIOS, firmware, and driver updates. If required, download the latest BIOS, firmware, and driver updates from the support site.
- Create a device driver media using one of the following methods:
 - USB drive
 - CD
 - DVD

Install PERC H965 series card driver during a Windows Server 2025 installation

About this task

To install the driver:

Steps

1. Start the system using the Windows Server 2025 media.
2. Follow the on-screen instructions until you reach **Where do you want to install Windows Server 2025** window, and then select **Load driver**.
3. When prompted to insert the media, insert the installation media and browse to the appropriate location.
4. Select **Controller HBA** from the list.
5. Click **Next** and continue the installation process.

Install PERC H965i Adapter, PERC H965i Front, PERC H965i MX, PERC H965e Adapter while newly installing Windows

About this task

To install the driver:

Steps

1. Boot the system using Windows or newer media.
2. Follow the on-screen instructions to navigate to the **Where do you want to install Windows** window, and then select **Load driver**.
3. When prompted, insert the installation media and browse to the appropriate location.
4. Select the driver files.
5. Click **Next** and continue installation.

Install PERC H965i Adapter, PERC H965i Front, PERC H965i MX, and PERC H965e Adapter driver on which Windows is already installed

About this task

Perform the following steps to configure the driver for the controller on which Windows is already installed:

Steps

1. Power off the system.
2. Install the new SAS controller in the system.
For instructions on installing the controller in the system, see [Install and remove the PERC 12 cards](#).
3. Power on the system.
The **Found New Hardware Wizard** screen displays the detected hardware device.
4. Click **Next**.
5. On the **Locate device driver** screen, select **Search for a suitable driver for my device** and click **Next**.
6. Browse and select the drivers on the **Locate Driver Files** screen.
7. Click **Next**.
8. Click **Finish**.
9. Reboot the system when prompted.

Update PERC H965i Adapter, PERC H965i Front, PERC H965i MX, and PERC H965e Adapter driver that runs on Microsoft Windows operating system

Prerequisites

 **NOTE:** Close all applications on your system before you update the driver.

Steps

1. Insert the media containing the driver.
2. Select **Start > Settings > Control Panel > System**.
The **System Properties** screen is displayed.

NOTE: The path to **System** may vary based on the operating system family.

3. Click the **Hardware** tab.

4. Click **Device Manager**.

The **Device Manager** screen is displayed.

NOTE: The path to **Device Manager** may vary based on the operating system family.

5. Expand **Storage Controllers** by double-clicking the entry or by clicking the plus (+) symbol next to **Storage Controllers**.

6. Double-click the controller for which you want to update the driver.

7. Click the **Driver** tab and click **Update Driver**.

The screen to update the device driver wizard is displayed.

8. Select **Install from a list or specific location**.

9. Click **Next**.

10. Follow the steps in the wizard and browse to the location of the driver files.

11. Select the INF file from the drive media.

12. Click **Next** and continue the installation steps in the wizard.

13. Click **Finish** to exit the wizard and reboot the system for the changes to become effective.

NOTE: Dell provides the Dell Update Package (DUP) to update drivers on systems running RHEL 8 or RHEL 9. DUP is an executable application that updates drivers for specific devices. DUP supports CLI and silent execution.

Linux driver installation

The Driver Update Disk (DUD) image files are created only for those operating system releases in which the native (in-box) driver is insufficient for installation. If an operating system is being installed with a corresponding DUD image file, see [Load the driver while installing an operating system](#). If not, use the native device driver and then:

- For RHEL 8 and RHEL 9: [Install or update a RPM driver package using the KMOD support](#).
- For SLES15 SP5: [Install or update a RPM driver package using the KMP support](#).

NOTE: To view the complete list of boot loader options, see the installation guide of your operating system.

Install or update a RPM driver package using the KMOD support

Prerequisites

NOTE: This procedure is applicable for RHEL 8.0 and later versions.

About this task

To install the RPM package with KMOD support, do the following:

Steps

1. Expand tarball driver release package.

2. Install the driver package by running the command: `rpm -ihv kmod-mpi3mr-<version>.rpm`.

NOTE: Use `rpm -Uvh <package name>` when upgrading an existing package.

3. If the previous device driver is in use, you must restart the system for the updated driver to take effect.

4. Verify the loaded driver version by running the following command: `modinfo mpi3mr`.

Install or update a RPM driver package using the KMP support

Prerequisites

 **NOTE:** This procedure is applicable for SUSE Enterprise Linux 15.x.

About this task

To install the RPM package with KMP support, do the following:

Steps

1. Expand the tarball driver release package.
2. Install the driver package by running the commands: `rpm -ivh broadcom-mpi3mr-kmp-<version>.rpm`

 **NOTE:** Use `rpm -Uvh broadcom-mpi3mr-kmp-<version>.rpm` to update an existing package.

3. If the previous device driver is in use, you must restart the system for the updated driver to take effect.
4. Verify the loaded driver version by running the following command: `modinfo mpi3mr`.

Load the driver while installing an operating system

About this task

 **NOTE:** Steps provided are general steps for the Linux based operating systems. For exact information about loading drivers while installing operating systems, see the Dell technical documentation of the operating system on the Dell support site or the official support of the operation system vendor.

Steps

1. Perform the following operation to install the driver media:
 - a. Download the PERC Linux driver ISO image file from the [Dell Support Site](#), or install the Lifecycle Controller driver pack.
 - b. Mount the ISO file to the server, burn the ISO to CD/DVD, or copy the ISO image file to USB. The USB has to match with the ISO file.
 - c. For Lifecycle Controller driver pack, boot the Lifecycle Controller and complete the tasks prompted by the operating system deployment wizard.
2. Boot to the installer.
3. On the Bootloader screen, select **E** and press Enter.
4. Do one of the following:

- If the operating system is RHEL, the CLI displays the syntax `vmlinuz`. Enter **inst.dd**.

For example, when you are prompted with the command `vmlinuz intrd=initrd.img inst.stage2=hd:LABEL=RHEL-7.0\x20x86_64 quiet inst.dd`.

- If the operating system is SLES, the CLI displays the syntax `linuxefi`. Enter **dud=1**.

For example, when you are prompted with the command `linuxefi/boot/x86_64/loader/linux splash=silent dud=1`.

 **NOTE:** Boot parameters may vary based on the operating system version. See operating system installation manuals for exact boot parameter syntax.

5. Attach the driver media (ISO, USB).
6. Press F10 to boot to the operating system.
7. Wait for the OS image file to boot and OS installer to begin. When prompted, select the driver media (for example, USB, CD, or an ISO image file).
8. When prompted, select the driver media.
If applicable, select the PERC driver `mpi3mr`.

 **NOTE:** Ensure that the selected driver is indicated by an **X** symbol.

9. Extract or load the driver.
10. Before proceeding or closing the **Driver Select** menu, disconnect the driver media.
 -  **NOTE:** Ensure that you disconnect the driver media so that the drivers are loaded successfully. If the installation media is deleted, reattach it.
11. Press C or select **Exit** to go to the installation.

Manage the PERC 12 firmware

This section provides information about downloading, installing, and upgrading the PERC 12 firmware using the Dell Update Package (DUP). The 17G PowerEdge servers support only PERC12 versions 8.11.0.0.18-22 and later. You cannot downgrade to an earlier version of PERC12 on 17G servers.

 **NOTE:** The PCIe Switch-based configuration support is available in firmware 8.4.0.0.18-27 and later versions.

Topics:

- [Upgrade firmware controller using Dell Update Package \(DUP\)](#)

Upgrade firmware controller using Dell Update Package (DUP)

About this task

 **NOTE:** If the Online Capacity Expansion operation is in progress then you cannot update the firmware version.

Steps

1. Go to the Drivers and Downloads page on the support site.
 2. Locate your controller.
 3. Download the DUP file.
 - a. To upgrade by using Windows or iDRAC, download the Windows executable file.
 - b. To upgrade using Linux, download the **.bin** file.
-  **NOTE:** For VMware, firmware must be upgraded by using iDRAC or the PERC CLI.
4. Install the DUP by doing one of the following:
 - a. For Windows, run the executable file in the Windows environment.
 - b. For Linux, run the **.bin** file in the Linux environment.
 - c. For iDRAC, click **System iDRAC > Maintenance > System Update**, upload Windows executable, and then install.

Manage PERC 12 controllers using HII Configuration Utility

The Human Interface Infrastructure (HII) configuration utility is a storage management application integrated into the System BIOS <F2>. It is used to configure and manage the controller(s), virtual disks, and physical disks. This utility is independent of the operating system.

Topics:

- [Enter the PERC 12 HII Configuration Utility](#)
- [Exit the PERC 12 HII Configuration Utility](#)
- [Navigate to the PERC 12 Configuration Utility](#)
- [View the HII Configuration Utility dashboard](#)
- [Configuration management](#)
- [Controller management](#)
- [Virtual disk management](#)
- [Device management](#)
- [Hardware components](#)
- [Security key management in HII configuration utility](#)

Enter the PERC 12 HII Configuration Utility

About this task

Perform the following steps to boot to the HII Configuration Utility:

Steps

1. Power on the server.
2. When the server starts, press F2 to open the **System Setup** page.
3. Click **Device Settings**.

The **Device Settings** page lists all the RAID controllers in the system.

To access the management menu of the controller, use the arrow keys or the mouse device.

 **NOTE:** For more information about each option, click Help in the upper-right corner. For information about individual options pause the pointer over each link and view the description in the bottom pane.

 **NOTE:** Some of the options within the HII Configuration Utility are not present if the controller does not support the corresponding feature. Options may also be grayed out if the feature is not applicable to the current configuration.

Exit the PERC 12 HII Configuration Utility

About this task

To exit the HII Configuration Utility, do the following:

Steps

1. Click **Finish** at the bottom-right corner on the **System Setup Main Menu** page.
2. When prompted to confirm if you want to perform the operation, click **Yes**.

Navigate to the PERC 12 Configuration Utility

Steps

1. Enter the UEFI configuration Utility. See [Enter the PERC 12 HII Configuration Utility](#). The **Device Settings** screen displays a list of NIC ports and the RAID controllers.
2. To enter PERC 12 configuration utility, click the appropriate PERC controllers. The **Dashboard view** screen is displayed.

View the HII Configuration Utility dashboard

The following table lists the features displayed in the Dashboard View page of the HII Configuration Utility:

Table 7. Dashboard view screen

Dashboard view options	Description
Main menu	Displays the following configuration options: <ul style="list-style-type: none">● Configuration Management● Controller Management● Virtual Disk Management● Device Management● Energy Pack Management
Help	Press F1 to view context-sensitive Help information.
Configuration Management	Displays the following options under Configuration Management : <ul style="list-style-type: none">● Create Virtual Disk—Creates a virtual disk by selecting RAID level, hard drives, and virtual disk parameters.● Create Profile Based Virtual Disk—Creates a virtual disk by using a wizard. This wizard makes intelligent choices based on the profile that you selected.● View Disk Group Properties—Displays associated virtual disks for the disk group and any available free capacity.● View Global Hot Spares—Displays the drives that are assigned as global hot-spare devices.● Convert to Non-RAID disk—Allows changing the state of the RAID capable disks to Non-RAID disks.● Clear Configuration—Deletes all existing configurations on the RAID controller.
Controller Management	Displays controller status and basic properties of the controller: <ul style="list-style-type: none">● Advanced Controller Management—Provides to links to various controller management activities.● Advanced Controller Properties—Displays controller cache and memory-related properties
Virtual Disk Management	Displays the properties of a specific virtual disk. You can perform operations such as Initialization and Check Consistency.
Device Management	Displays logical enclosure details and the physical disks that are attached to it.
Energy Pack Management	Displays the details such as its type, status, capacity, and so on about the energy pack.

Configuration management

Create virtual disk and configure virtual disk parameters

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See, [Navigating to the Dell PERC 12 configuration utility](#).

2. Click **Main Menu > Configuration Management > Create Virtual Disk**.
3. Select the RAID level. You can select PDs from either unconfigured capacity or free capacity. The list of physical disks appear.
4. Click **Select Physical Disks**. See [Select physical disks for creating VDs](#).
5. Select the physical disks for the selected RAID level and click **OK**.
6. Click **Confirm** and click **Yes**.
7. Click **OK**. The **CONFIGURE VIRTUAL DISK PARAMETERS** section is displayed.

 **NOTE:** Mixing of TCG Enterprise and TCG Opal SED protocols in a virtual disk is not supported.

Table 8. Virtual disk parameters and their descriptions

Virtual disk parameters	Description
Virtual Disk Name	Enter the name of the virtual disk.
Virtual Disk Size	Displays the maximum capacity available for the virtual disk.
Virtual Disk Size Unit	Displays the virtual disk storage space either in GiB or TiB.
Strip Element Size	Select the strip element size. The disk striping involves partitioning each physical disk storage space in stripes of the sizes 64 KiB and 256 KiB.
Read Cache Policy	Displays the controller read policy. By default, the read cache policy is set to No Read-Ahead . You can set the Read policy to: <ul style="list-style-type: none"> ● Read-Ahead—The controller speeds up sequential Read operations on Write-Back virtual disks. This is supported only on HDDs. ● No Read-Ahead— Read-Ahead is not enabled for the virtual drive
Write Cache Policy	Displays the controller write cache policy You can set the write policy to: <ul style="list-style-type: none"> ● Write-Through—The controller sends a data transfer completion signal to the host when the disk subsystem has received all the data in a transaction. ● Write-Back—The controller sends a data transfer completion signal to the host when the controller cache has received all the data in a transaction. ● Force Write-Back—The controller sends a data transfer completion signal to the host when the controller cache has received all the data in a transaction, even if the controller energy pack is below the threshold charge. By default, the Write Cache Policy is set to Write-Back .
Disk Write Cache Policy	Select the disk cache policy to Default , Enable , or Disable . By default, the Disk Write Cache Policy is set to Default .
Default Initialization	Select the default initialization to: <ul style="list-style-type: none"> ● No —Virtual disk is not initialized. ● Fast—The first and last 8 MB of the virtual disk is initialized. ● Full—Entire virtual disk is initialized. For more information, see Virtual disk initialization . By default, the default initialization is set to No.

8. Click **Create virtual disk**.
9. Click **Confirm** and click **Yes**. The virtual disk is created.

Select physical disks for creating VDs

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See, [Navigating to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Configuration Management > Create Virtual Disk**.
3. Click **Select Physical Disks**.
4. Select the media type. For example, SSD, HDD, or both. Based on your selection, the drives are displayed for creating RAID drives.
5. Select the interface type. For example, SAS, SATA, or NVMe.

6. Select the logical sector size. For example, 512B, 4KiB, or both.

Create profile based virtual disk

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigating to the Dell PERC 12 configuration utility](#) .
2. Click **Main Menu > Configuration Management > Creating Profile Based Virtual Disk**.
The following list of RAID modes are displayed:

- **Generic RAID 0**
- **Generic RAID 1**
- **Generic RAID 5**
- **Generic RAID 6**
- **File Server**
- **Web/Generic Server**
- **Database**

Based on the RAID mode selected, one or more the physical disk selection criteria is displayed.

3. From the **Physical Disk Selection Criteria** drop-down menu, select a criterion based your requirement.
The Profile Parameters of the selected option is displayed.

i **NOTE:** Based on the physical disk selection criteria, if enough disks are not available to satisfy the applicable RAID level, then the **Physical Disk Selection Criteria** option is grayed out.

i **NOTE:** Before using drives for profile based VD creation, ensure that the drives are securely erased. When selecting a drive, ensure that you do not mix:

- HDDs and SSDs.
- NVMe, SAS, and SATA types.
- Drives that support only Physical Region Page (PRP) and that support both PRP and Scatter Gather List (SGL).
- 512b and 4K block-size drives.
- Drives of different link speeds such as 3G, 6G, 12G, and 24G, or 2.5GT, 5.0GT, 8.0GT, 16.0GT, or 32.0GT.
- Drives of SED and non-SED types.
- SEDs are not secured.
- Drives of single and multiple LUN types.

4. Click **Create Virtual Disk**.
5. Select **Confirm** and click **Yes** to continue.
The virtual disk is created with the parameters of the profile selected.

View disk group properties

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigating to the Dell PERC 12 configuration utility](#) .
2. Click **Main Menu > Configuration Management > View Disk Group Properties**.
The list of disk group properties are displayed:

Table 9. disk group properties

Properties	Descriptions
Capacity Allocation	Displays all the virtual disks associated with the specific disk group. It also provides information about the available free disk space.
Secured	Displays whether the disk group is secured or not.

View global hot-spare devices

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigating to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Configuration Management > View Global Hot-Spares**.
The list of PDs that are allocated as global hot-spares is displayed.

Convert to Non-RAID disk

To convert a physical disk to non-RAID disk using the HII Configuration Utility, do the following:

Steps

1. Start the Dell PERC 12 Configuration Utility. See [Navigating to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Configuration Management > Convert to Non-RAID Disk**.
A list of available PDs is displayed.
3. Select the physical disk to convert to Non-RAID disk.
4. Click **OK**.
5. Click **Confirm**.
6. Click **Yes**.
The non-RAID disk is successfully created.

Delete configurations

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigating to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Configuration Management > Clear Configuration**.
3.  **CAUTION: It is recommended that you back up data stored on the virtual disks and hot-spare disks on the controller before deleting the virtual drive.**

When prompted to indicate if you want to delete VD's and hot-spares, select the **Confirm** check box, and then click **Yes**. The virtual disks and hot-spare disks available on the controller are deleted successfully.

Controller management

Save persistent events

To save persistent events, you must have a USB or file system on the server.

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Management**.
3. Click **Save Persistent Events**.
4. When prompted to indicate if you want to save persistent events, click **OK**.

Clear persistent events

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).

2. Click **Main Menu > Controller Management > Advanced Controller Management**.
3. Click **Clear Persistent Events**.
4. Click **OK**.

Manage snapdump

Prerequisites

To save the snapdump, you must have a USB or a file system attached to the server before starting the server. For information about the Snapdump feature, see [Snapdump](#).

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Management**.
3. Click **Manage Snapdump**.
4. In the OPERATIONS section, select a file system and its corresponding directory. To save the Snapdump data of the selected file, click **Save new On-Demand Snapdump**. To delete the Snapdump data, click **Clear All Snapdumps**.
5. Click **OK**.

Enable security

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigating to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Management**.
3. Click **Enable security**, select **Local Key Management**.
4. Click **Ok**.
5. If you want to use the security key generated by the controller, click **Suggest Security Key** and **OK**.
The operation is successful.
6. Enter the key in the **Security Key** and **Confirm** boxes.
7. Select the **I Recorded the Security Settings For Future Reference** check box and click **Enable Security**.
8. Select **Confirm** and click **Yes**.
9. Click **OK**.
The operation is successful.

Disable security

You can disable the LKM security feature by using the HII.

About this task

 **NOTE:** Before disabling the drive security feature, ensure that all secured drives are either erased or removed.

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the PERC 12 Configuration Utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Management**.
3. Click **Disable security**.
4. When prompted if you want to disable security, click **Confirm**, and then click **Yes**.

Change security settings

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Management**.
3. Click **Change Security Settings**, select **Change Current Security Settings**.
4. Click **Ok**.
5. If you want to use the security key generated by the controller, click **Suggest Security Key** and confirm the security key by re-entering.
The operation is successful.
6. Click **Save Security Settings**.
7. Select **Confirm** and click **Yes**.
8. Click **Ok**.

Restore factory default settings

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Management > Set Factory Defaults**.
3. Select **Confirm** and click **Yes**.

Manage the SAS Storage link speed

About this task

To change the storage link speed of the eligible SAS/SATA disks, do the following:

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Management > Manage SAS Storage Link Speed**.
3. To change the link speed, select the appropriate options for each link, and then click **Apply Changes**.
4. Click **OK**.

Manage the PCIe storage interface

About this task

To change the storage interface of the eligible NVMe disks, do the following:

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Management > Manage PCIe Storage Interface**.
3. To change the link speed, click **View/Change Link Speed** and select the appropriate options for each port, and then click **Apply Changes**.
4. To manage the PCIe width, select the appropriate width for each available link, and then click **Apply Changes**.
5. Select the **Confirm** check box and click **Yes**.
6. Click **OK**.

 **NOTE:** Physical drive slot numbers do not correlate with the controller PHY numbers. The physical drive slot numbers displayed on the HII may not be the same slot numbers in which the drives are actually installed.

Auto-Configure Behavior Management

About this task

To change the Auto-Configure Behavior of the eligible disks, do the following:

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Management > Manage Controller Personality**.
3. To change the Auto-Configure behavior (Primary or Secondary or both) on all the new disks that will be inserted, select the appropriate options for **Auto-Configure Behavior (Primary)** and **Auto-Configure Behavior (Secondary)**, and then click **Apply Changes**.
4. To run immediate one-time Auto-Configure operation, select the **Auto-Configure Behavior (Execute Once)** option, and then click **Apply Now**.

Advanced controller properties

Set the patrol read mode

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Properties**.
3. Click **Patrol Read**.
The following options are displayed:
 - **Start**—Starts patrol-read for the selected controller.
 - **Suspend**—Suspends the ongoing patrol-read operation on the controller.
 - **Resume**—Resumes the suspended patrol-read operation.
 - **Stop**—Stops patrol-read for the selected controller.
 - **Rate**—Indicates the percentage of system resources dedicated to perform the patrol-read operation.
4. To enable or disable patrol-read and to correct unconfigured areas on the physical disk, click **Correct Unconfigured Areas**.
5. Click **Apply Changes**.

Configure hot-spare drives

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Properties**.
3. Click **Spare**.
The following options are displayed:
 - **Replace Member**—Enables or disables the option to replace the member.
 - **Auto Replace Member on HDD Predictive Failure**—Enables or disables the option to start a Replace Member operation if a predictive failure error is detected on a PD.
 - **Auto Replace Member on SSD Predictive Failure**—Enables or disables the option to start a Replace Member operation if a predictive failure error is detected on an SSD.
4. Select the applicable option and click **Apply Changes**.
The changes are saved successfully.

Set task rates

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Properties**.

3. Click **Task Rates**.

The following options are displayed:

- **Background Initialization (BGI) Rate**
- **Patrol Read Rate**
- **Consistency Check Rate**
- **Rebuild Operating Mode Priority:** Available in 8.8.0.0.18-26 and later versions of PERC12 controllers.
 - Rebuild—The rebuild operation is given priority over the host system's IOPS.
 - Host I/O—The host Input/Output operations are given priority over the rebuild operation.
- **Online Capacity Expansion (OCE)**

4. You can make the necessary changes and then **Apply Changes**.

The task rates operation is completed successfully.

Controller properties

Set a drive as the first device to OS

About this task

The First Device feature is available only on Linux based servers. Using the First Device feature, you can select the order in which the PD or VD must be presented to the OS. You can also select the PD or VD on which you can install the OS Do the following:

Steps

1. [Enter the PERC 12 HII Configuration Utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Properties**.
3. In the **Controller Properties** section, from the **First Device** drop-down menu, select the PD or VD that must be set as the first device.
4. Click **Apply Changes**.
The selected Physical Device or Virtual Drive is set as the first device.

Device reporting order

About this task

The Device Reporting Order feature is available only on Linux based servers. Enables you select the order in which either non-RAID drives or VDs must be listed first to the OS. Do the following:

Steps

1. [Enter the PERC 12 HII Configuration Utility](#)
2. Click **Main Menu > Controller Management > Advanced Controller Properties**.
3. In the **Controller Properties** section, from the **Device Reportig Order** drop-down menu, select to either first present a non-RAID drive or VD to the OS.
4. Click **Apply Changes**.
The devices are presented to the OS based on the order your selected.

Manage the write-cache feature of NVMe devices

About this task

To enable the write-cache feature on NVMe devices, do the following:

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Properties**.
3. In the **Controller Properties** section, set the **Write Cache for NVMe Devices** feature to Enabled, Disabled, or Default.

4. Click **Apply Changes**.
The write-cache policy of NVMe devices is updated.

Manage the write-cache feature of SATA devices

About this task

To enable the write-cache feature on SATA devices, do the following:

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Properties**.
3. In the **Controller Properties** section, set the **Write Cache for SATA Devices** feature to Enabled, Disabled, or Default.
4. Click **Apply Changes**.
The write-cache policy of SATA devices is updated.

Select boot mode

About this task

To select the boot mode, perform the following steps:

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Properties**.
3. In the **Controller Properties** section, select either **Continue on Errors** or **Safe Mode on Errors** from the **Boot Mode** list.

 **NOTE:** By default, the boot mode option is set to **Continue on Errors**.

Table 10. Boot mode options

Option	Description
Continue on Errors	The controller attempts to automatically clear errors and continue booting. Errors can result in safe mode if the controller is unable to clear them.
Safe Mode on Errors	The controller is routed to safe mode when critical errors arise. PERC firmware disables most of the features on the controller. The controller requires attention from the user to acknowledge and rectify the issues.

4. Click **Apply Changes**.
The boot mode operation is completed successfully.

Abort the consistency check

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Properties**.
3. In the **Controller Properties** section, set the **Abort Consistency Check on Error** option to **Enabled** or **Disabled**.
4. Click **Apply Changes**.
The option to abort the consistency check operation on a redundant virtual disk is enabled if there is any inconsistency found in the data.

Discard preserved cache

By using the Discard Preserved Cache feature, you can delete the preserved cache on the controller that could not be written to a VD after the power is restored and the Write Cache operation is restarted.

About this task

 **WARNING:** Use this option only if necessary because it may result in data loss. Before discarding the preserved cache, it is recommended to use all available data recovery techniques and correct the issue.

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the PERC 12 Configuration Utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Properties**.
3. Click **Cache and Memory > Discard Preserved Cache**.
The preserved cache is cleared successfully.

Virtual disk management

Virtual disk numbering

Virtual disks are numbered in ascending order beginning from the lowest, which is ID 1.

View virtual disk properties

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Enter the PERC 12 HII Configuration Utility](#).
2. Click **Main Menu > Virtual Disk Management**.
All the virtual disks that are associated with the RAID controller are displayed.
3. To view the properties, click the virtual disk. You can view the following properties of the virtual disk:

Table 11. Virtual disk properties

Option	Description
Operation	Select one of the following: <ul style="list-style-type: none">● Delete Virtual Disk: Deletes the selected VDs.● Fast Initialization● Full Initialization: See Full initialization.● Consistency Check (available for all RAID levels except for RAID 0 VD). See Check consistency of VDs.● Expand Virtual Disk
Name	Indicates the name of the virtual disk.
RAID level	Indicates the RAID level of the virtual disk.
Status	Indicates the status of the virtual disk. The possible options are: <ul style="list-style-type: none">● Ready● Degraded● Offline● Failed
Size	Indicates the size of the virtual disk.
Disk Group	Indicates the disk group of the virtual drive.

View physical disks associated with a virtual disk

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Enter the PERC 12 configuration utility](#).
2. Click **Main Menu > Virtual Disk Management**.
All the virtual disks associated with the RAID controller are displayed.
3. Click on a virtual disk.
The properties of the virtual disk are displayed.
4. Click **View Associated Physical Disks**.
All the physical disks that are associated with the virtual disk are displayed.
5. From the **Associated Physical Disks** section, select the physical disk.
6. Click **View Physical Disk Properties** to view the physical disk properties.

View advanced properties of a virtual disk

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Enter the PERC 12 HII Configuration Utility](#).
2. Click **Main Menu > Virtual Disk Management**.
All the virtual disks that are associated with the RAID controller are displayed.
3. Click the virtual disk.
The properties of the virtual disk are displayed.
4. Click **Advanced**.
You can view the following additional properties of the virtual disk:

Table 12. Advanced properties of the virtual disk

Option	Description
Strip element size	Indicates the strip element size for the virtual disk.
Secured	Indicates whether the virtual disk is secured or not.
Bad blocks	Indicates whether the virtual disk has corrupted blocks.
Data format for I/O	Indicates the data format for I/O operations (only available for NVMe drives)

Configure virtual disk policies

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Enter the PERC 12 configuration utility](#).
2. Click **Main Menu > Virtual Disk Management**.
All the virtual disks associated with the RAID controller are displayed.
3. Select the virtual disk.
4. Click **Advanced**.
You can view the following virtual disk policies:

Table 13. Virtual disk policies

Option	Description
Current Cache Status	Indicates the current read and write cache policies for the virtual disk.
Access Policy	Indicates the current access policy for the virtual disk.
Default Read Cache Policy	Allows selection of the read cache policy for the virtual disk. The available option is Read Ahead .
Current Power Save Policy	Indicates the current power saving policy of the VD.

Table 13. Virtual disk policies (continued)

Option	Description
Default Write Cache Policy	Allows selection of the write cache policy for the virtual disk. The available options are: Write Through , Write Back , and Force Write Back .
Disk Write Cache Policy	Allows selection of the Disk Write Cache Policy for the virtual disk. The possible options are: <ul style="list-style-type: none"> • Default (Disk Default) • Enable • Disable

5. Click **Apply Changes**.
The changes made are saved successfully.

Configure Virtual Disks

When configuring the virtual disks, you should consider the workload intended. For example, RAID 1 for a simple boot disk; RAID 5 or RAID 6 for file or web servers (sequential reads/writes of files); or RAID 10 for transactional database (small random-read and random-write operations). Virtual disks configured on hard drives should use the cache settings of write-back and read-ahead.

Virtual disks configured on SSDs can use the same controller default settings as used by hard drives. Most users copy data of OS files or database to the new array. Using such default settings provide optimum performance in this configuration. After copying, the array can be used as-is based on the number and type of SSDs.

Perform Online capacity expansion—Expand a VD within a disk group

About this task

Expand a VD within a disk group

 **NOTE:** This feature is available only if there is sufficient free disk space available in the associated disk group.

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Enter the PERC 12 configuration utility](#).
2. Click **Main Menu > Virtual Disk Management**.
3. Select the virtual disk.
4. From the **Operations** drop-down menu, select **Expand Virtual Disk**.
 **NOTE:** The expand virtual disk feature is available only if there is free space available in the associated disk group
5. Click **Go**.
6. To expand a VD, enter the percentage of available capacity and click **Expand**. The VD is successfully expanded.

Perform Online capacity expansion—Expand a VD by adding a PD to a disk group

About this task

Expand a VD by adding physical drive to a disk group

 **NOTE:** This feature is available only if there is no sufficient free disk space available in the associated disk group.

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Enter the PERC 12 configuration utility](#).

2. Click **Main Menu > Virtual Disk Management**.
3. Select the virtual disk.
4. From the **Operations** drop-down menu, select **Expand Virtual Disk**.
5. Click **Go**.
6. Click **Add Physical Disks**.
7. To expand a VD, select physical disk(s) that you want to add to the VD.
8. Click **OK**. The VD is successfully expanded.

Check consistency of VDs

Prerequisites

To enable consistency check from the HII Configuration Utility, do the following:

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Virtual Disk Management**.
The list of virtual disks is displayed.
3. Select the virtual disk.
 **NOTE:** Consistency check cannot be run on RAID 0 virtual disks.
4. From the **Operations** list, select **Check Consistency**.
5. Click **Go**.
6. When prompted to indicate if you want to start the checking operation, select the **Confirm** check box, and then click **Yes**.

Device management

View enclosure or backplane properties

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Device Management > Logical Enclosure <enclosure number>**. An enclosure number can be 0, 1, 2, 3, 4, and so on.
All the physical disks associated with the selected enclosure are listed.
 **NOTE:** Fields associated with properties that are not applicable are hidden.

Table 14. Enclosure or backplane properties

Option	Description
Enclosure ID	Displays the persistent ID of the enclosure.
Bay ID	Displays the Bay ID of the enclosure.
Name	Displays the name of the enclosure.
Connector name	Indicates the type of physical connection to the enclosure. All connectors for internal controllers are identified by "00".
Position	Position of an enclosure in the enclosure stack.
State	Displays the state of the enclosure.
Vendor ID	Displays the vendor ID of the enclosure.
Product ID	Displays the product ID of the enclosure.

Table 14. Enclosure or backplane properties (continued)

Option	Description
Location	Displays the location of the enclosure. Location is either internal or external.
Type	Displays the type of the enclosure. The types of enclosures are Virtual SES , SAS Expander , and Managed PCIe Switch .
Product Revision Level	Displays the product revision level of the enclosure.
Number of Slots	Displays the number of slots in the enclosure.
Number of Physical Disks	Indicates the number of PDs installed on the enclosure.
Partner Connector Name	Indicates the connector name of the partner device.
Partner Position	Indicates the position of the partner device.
Partner Enclosure ID	Indicates the enclosure ID of the partner.

 **NOTE:** "Partner" refers to the devices that have multiple paths to communicate for redundancy and load balancing.

 **NOTE:** The Partner entries are displayed only for backplane devices in multipath. Else, it will not be displayed.

View physical disk properties

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Device Management > Logical Enclosure <enclosure number>**. An enclosure number can be 0, 1, 2, 3, 4, and so on.
The enclosure properties and all the physical disks associated with the selected enclosure are listed.
3. To view the properties of a PD, click the corresponding PD link.

 **NOTE:** Fields associated with features that are not applicable are hidden.

Table 15. Physical disk Dashboard View

Option	Description
Operation	The list of operations you can perform on the selected physical disk. The options are: <ul style="list-style-type: none"> • Blink • Unblink • Assign global hot spare • Cryptographic erase • Convert to non-RAID disk • Make Offline • Replace Member
Slot number	The PD slot to which the controller is connected.
Persistent ID	Persistent ID of the physical drive.
Status	Status of the physical disk.
Size	Size of the physical disk.
Type	Type of the physical disk.
Model	Model of the physical disk.
Part Number	Part number of the physical drive.

Table 15. Physical disk Dashboard View (continued)

Option	Description
Serial Number	Serial of the physical disk.
Vendor	Name of the PD manufacturer.
Firmware Revision Level	Indicates the version of the firmware installed on the PD.
Manufacturing Date	Date of manufacture of the PD.
Associated Virtual Disk	Indicates the VD that is associated with this PD.

4. To view additional properties of the physical disk, click **Advanced**.

Table 16. Advanced physical disk properties

Option	Description
SMART status	SMART status of a physical disk
SAS Address/WWID	WWN number of the device.
Disk Power Status	Power condition (On or Power Save) of the physical disk
Interface	Indicates the type of device used as an interface. For example, SAS or parallel SCSI.
Capable Speed	Indicates the read-write speed capability of PD in Gbps.
Negotiated Speed	Negotiated link speed of the device.
Capable Link Width	Capable link width of the device.
Negotiated Link Width	Negotiated link width of the device.
Number of Connections	Indicates the number of paths do the PD.
Cryptographic erase capable	Cryptographic erase capability of the physical disk
Encryption Capable	Indicates whether or not the drive can be encrypted.
Supported Data Format	Indicated for NVMe drives only.
Temperature (C)	Indicates the current temperature of the PD.

5. To view information about the Logical Unit Number (LUN) and NVMe Namespace properties of a PD, click **Logical Unit/ Namespace Information**.

Table 17. LUN/Namespace data of a physical disk

Option	Description
NVMe Namespace ID	Indicates LUN and Namespace properties
Status	Indicates the working status of a LUN.
Size	Indicates the maximum storage size of the LUN.
Total Unconfigured Space	Total free space available for RAID array.
Total Configured Space	Total used drive space for RAID array.
Logical Sector Size	Supports either 512B or 4KiB types.
Physical Sector Size	Supports either 512B or 4KiB types.
Media Error Count	Number of physical errors detected on the PD.
Other Error Count	Other errors detected on the PD.
Predictive Failure Count	Predictive errors detected on the PD.
Firmware Managed Security	Indicates whether or not the PD security is managed by the controller firmware.
Current Write-Cache	Indicates whether or not the current write-cache mode of the PD is enabled.

Table 17. LUN/Namespace data of a physical disk (continued)

Option	Description
Default Write-Cache	The default write-cache mode of the PD.
Secured	Indicates whether or not the PD is secured.
Locked	Indicates whether or not the PD is locked.

To view information about the next PD associated with the enclosure, click **New Physical Disk**.

Cryptographic erase

Cryptographic erase is a process to erase all data permanently on an encryption-capable and unconfigured physical disk, and reset the security attributes. Cryptographic erase on an Self Encrypting Drives (SED) will unsecure the disk.

Prerequisites

- The non-RAID and virtual disks associated with the drive are deleted.
- The disks are not dedicated as hot-spares.

About this task

The Cryptographic erase feature is supported only on Instant Secure Erase (ISE) and SEDs.

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Device Management > Logical Enclosure <enclosure Number>**. An enclosure number can be 0, 1, 2, 3, 4, and so on.
All the physical disks associated with the selected enclosure are listed.
3. Select a physical disk.
4. From the **Operations** list, select **Cryptographic Erase**.

 **NOTE:** The Cryptographic Erase option is displayed only if the drive installed is ISE or SED capable.

5. Click **Go**.
6. When prompted to indicate if you want to start the cryptographic erase operation, select the **Confirm** check box, and then click **Yes**.
The cryptographic erase operation is successfully completed.

Cryptographic Erase with PSID Revert

PSID erase is a process to erase all data permanently on an encryption-capable and unconfigured physical disk, and reset the security attributes. PSID erase on an Self Encrypting Drive (SED) will unsecure the disk.

Prerequisites

- The non-RAID drives associated with the drive are deleted.
- Remove the drive, note down the PSID printed on the drive so you can enter in the PSID field of HII, and then insert the drive back into the slot.

About this task

The PSID erase feature is supported only on SEDs.

Steps

1. Enter the **Dell PERC Configuration Utility**.
2. Click **Main Menu > Device Management > Logical Enclosure <enclosure Number>**. An enclosure number can be 0, 1, 2, 3, 4, and so on.

All the physical disks associated with the selected enclosure are listed.

3. Select a physical disk.
4. From the **Operations** list, select **Cryptographic Erase with PSID Revert**.
 **NOTE:** The Cryptographic Erase with PSID Revert option is displayed only if the drive installed is SED capable.
5. Enter the PSID key of the drive.
6. Click **Go**.
7. When prompted to indicate if you want to start the erase operation, select the **Confirm** check box, and then click **Yes**.
The Cryptographic Erase with PSID Revert operation is successfully completed.

Physical disk erase

Prerequisites

To use the Physical Disk Erase feature from the HII Configuration Utility, perform the following steps:

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Device Management > Logical Enclosure <enclosure Number>**. An enclosure number can be 0, 1, 2, 3, 4, and so on.
All the physical disks associated with the selected enclosure are listed.
3. Select a physical disk.
4. From the **Operations** drop-down menu, select **Physical Disk Erase**.
 **NOTE:** The Cryptographic Erase option is displayed only if the drive is of either ISE or SED type.
5. Click **Go**.
6. When prompted to indicate if you want to start erasing data on the physical disk, select the **Confirm** check box, and then click **Yes**.
Data on the physical disk is successfully erased.

Sanitize Block or Overwrite Erase operation

Sanitize secure erase is a process that a drive performs to erase all data permanently by modifying physical block to a user or vendor-specific value. Sanitize can be managed by the PERC or by an application running on the OS. See the Physical Disk Erase section in this guide.

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Device Management > Logical Enclosure <enclosure Number>**. An enclosure number can be 0, 1, 2, 3, 4, and so on.
All the physical disks associated with the selected enclosure are listed.
3. Select a physical disk. Ensure that the drive is not configured.
4. From the **Operations** list, select the required Sanitize Erase operation.
 **NOTE:** The Cryptographic Erase option is displayed only if the drive installed is ISE or SED capable.
5. Click **Go**.
6. When prompted to indicate if you want to start the sanitize erase operation, select the **Confirm** check box, and then click **Yes**.
The sanitize erase operation is successfully completed.

Assigning a global hot-spare

To assign a global hot spare from the HII Configuration Utility, do the following:

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Device Management > Logical Enclosure <enclosure Number>**. An enclosure number can be 0, 1, 2, 3, 4, and so on.
All the physical disks associated with the selected enclosure are listed.
3. Select the physical disk.
4. From the **Operations** list, select **Assign Global Hot Spare**.
5. Click **Go**.
6. When prompted to indicate if you want to assign the selected PDs as hot-spares, select the **Confirm** check box, and then click **Yes**.
The PDs are assigned as global hot-spare devices.

Assign a dedicated hot-spare

Prerequisites

To assign a dedicated hot spare from the HII Configuration Utility, perform the following steps:

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the Dell PERC 12 configuration utility](#).
2. Click **Main Menu > Device Management > Logical Enclosure <enclosure Number>**. An enclosure number can be 0, 1, 2, 3, 4, and so on.
All the physical disks associated with the selected enclosure are listed.
3. Select the physical disk.
4. From the **Operations** drop-down menu, select **Assign Dedicated Hot Spare**.
5. Click **Go**.
6. When prompted to indicate if you want to assign the selected PDs as dedicated hot-spare devices, select the **Confirm** check box, and then click **Yes**.

i **NOTE:** When all drives associated with a VD (including the dedicated hot-spare assigned to that VD) are removed from the server and reinserted, the dedicated hot-spare will be converted to a global hot-spare if there are other VDs available in the server. If there are no VDs left in the server then the dedicated hot-spare will be marked as foreign drives.

Convert a non-RAID disk to RAID-capable

To convert a non-RAID disk to RAID capable disk from the HII Configuration Utility, perform the following steps:

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Navigate to the PERC 12 Configuration Utility](#).
2. Click **Main Menu > Device Management > Logical Enclosure <enclosure Number>**. An enclosure number can be 0, 1, 2, 3, 4, and so on.
All the physical disks associated with the selected enclosure are listed.
3. Select the physical disk.
4. From the **Operations** list, select **Convert to RAID capable**.
5. Click **Go**.
A screen appears asking if you are sure you want to perform the operation.
6. Select the **Confirm** option.
7. Click **Yes**.
The operation is successful.

Convert a physical disk to a non-RAID disk

To convert a physical disk to non-RAID disk from the HII Configuration Utility, perform the following steps:

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Enter the PERC 12 HII Configuration Utility](#).
2. Click **Main Menu > Device Management > Logical Enclosure <enclosure Number>**. An enclosure number can be 0, 1, 2, 3, 4, and so on.
All the physical disks associated with the selected enclosure are listed.
3. Select the physical disk.
4. From the **Operations** list, select **Convert to non-RAID Capable**.
5. Click **Go**.
6. When prompted to indicate if you want to convert the selected PDs to non-RAID types, select the **Confirm** check box, and then click **Yes**.
The PDs are converted to non-RAID PDs.

Hardware components

View battery (Energy Pack) properties

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Enter the PERC 12 HII Configuration Utility](#).
2. Click **Main Menu > Energy Pack Management**.
The battery (Energy Pack) and capacity information are displayed.
3. You can view the following properties of the battery (Energy Pack):

Table 18. Battery (Energy Pack) properties

Field	Description
Type	Displays the type of battery.
Status	Displays the current status of the battery.
Manufacturer	Displays the manufacturer of the battery.
Module Version	Displays the version of the module.
Design Capacity (mAh)	Displays the size of the battery.
Temperature (C)	Displays the current temperature of the battery and also indicates whether the temperature is normal or high.
Voltage (mV)	Displays whether the voltage status of the battery is normal or high.

4. To set advanced properties click **Advanced**.

Table 19. Advanced battery (Energy Pack) properties

Field	Description
Full capacity	Displays the maximum charge capacity of the battery.
Remaining capacity	Displays the current charge capacity of the battery.
Charge %	Displays the available charge of the battery in percentage.
Current (mA)	Displays power consumption of the battery in milliamps (mA).
Charge Cycle Count	Displays the number of times the battery is charged or discharged.
Auto-learn mode	Displays the condition of the battery.

Table 19. Advanced battery (Energy Pack) properties (continued)

Field	Description
Next Learn Cycle Time	Displays the next schedule of the auto learn cycle time.

5. To start a manual learning process, click **Start Manual Learn Cycle**.

Security key management in HII configuration utility

The Dell OpenManage storage management application and the **HII Configuration Utility** of the controller allow security keys to be created and managed as well as create secured virtual disks. The following section describes the menu options specific to security key management and provides instructions to perform the configuration tasks. The contents in the following section apply to the **HII Configuration Utility**.

- The **Controller Management** screen displays controller information and action menus. You can perform the following security-related actions through the controller management menu:
 - **Security Key Management**—Create, update, or delete a Local Key Management (LKM) security key.
- The **Virtual Disk Management** screen displays physical disk information and action menus. You can perform the following security related actions through the virtual disk management menu:
 - **Secure Disk Group**—Secures all virtual disks in disk group.
 - **Create secure virtual disk**—Creates a new virtual disk that is secured with the security key on the controller.
- The **Device Management > Enclosure X** screen displays physical disk information and action menus. You can perform the following security-related actions through the physical disk management menu:
 - **Secure non-RAID disk**—Secures the non-RAID disk with the controller security key.
 - **Cryptographic Erase**—Permanently erases all data on the physical disk and resets the security attributes.

For more information on the Device Management screen and the Virtual Disk Management screen, see [Device management](#) and [Virtual disk management](#).

Security key and RAID management

Topics:

- [Security key implementation](#)
- [Local Key Management](#)
- [Create a security key](#)
- [Change security settings](#)
- [Disable security key](#)
- [Create a secured virtual disk](#)
- [Secure a non-RAID disk](#)
- [Secure a pre-existing virtual disk](#)
- [Import a secured non-RAID disk](#)
- [Import a secured virtual disk](#)
- [Dell Technologies OpenManage Secure Enterprise Key Manager \(SEKM\)](#)

Security key implementation

The PERC 12 series of cards support Self-Encrypting Drives (SEDs) for protection of data against loss or theft of SEDs. Protection is achieved by the use of encryption technology on the drives. There is one security key per controller. You can manage the security key using Local Key Management (LKM) or OpenManage Secure Enterprise Key Manager, also referred to as Secure Enterprise Key Manager (SEKM). The LKM key can be escrowed into a file using Dell OpenManage Storage Management application. The security key is used by the controller to lock and unlock access to encryption-capable physical disks. To use this feature, you must:

1. Have SEDs in your server.
2. Have created a security key.

 **NOTE:** In external enclosures and a C6600 PowerEdge server, if the main system is shutdown, drives will remain in unlocked state until the drives are power cycled.

 **NOTE:** The iDRAC Auto Secure feature does not secure drives associated with PERC12 controllers.

Local Key Management

You can use Local Key Management (LKM) to generate the key ID and the Security Key that is required to secure the physical disks. You can secure physical disks, change security keys, and manage secured foreign configurations using this security mode.

Create a security key

About this task

 **NOTE:** There is no security key backup option when you create a security key; you need to remember your security key.

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Enter the PERC 12 HII Configuration Utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Management > Enable Security**.
3. Select the **Security Key Management** mode as the local key management type.
4. Click **Ok**.

5. In the **Security Key Identifier** box, enter an identifier for your security key.

i **NOTE:** The Security Key Identifier is a clear-text label that enables you to associate the correct security key with the controller.

6. If you want to use the security key generated by the controller, click **Suggest Security Key**.
Assigns a security key suggested by the controller automatically.

7. In the **Security Key** box, enter the security key.

i **NOTE:** Security key is case-sensitive. You must enter a minimum of 8 or a maximum of 32 characters. Ensure that the characters contain at least one number, one lower case letter, one upper case letter, and one non-alphanumeric character.

8. In the **Confirm** box, re-enter the security key.

i **NOTE:** If the security key entered in the **Security Key** and **Confirm** boxes do not match, then you are prompted with a message to reenter the security key.

9. Select the **I Recorded the Security Settings for Future Reference** check box.

10. Click **Enable Security**.

The Security Key is created successfully.

Change security settings

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Enter the PERC 12 HII Configuration Utility](#).

2. Click **Main Menu > Controller Management > Advanced Controller Management > Change Security Settings**.

3. Do the following:

- To change the security key identifier, enter a new key identifier in the **Enter a New Security Key identifier** box.
- To keep existing key identifier, select the **Use the existing Security Key Identifier** check box.

4. Enter the existing security key.

5. Set security key:

- To change the security key, enter a new security key in the **Enter a New Security Key** box. Re-enter the new security key to confirm.
- To keep the existing security key, select the **Use the Existing Security Key** check box.

6. Select the **I recorded the Security Settings for Future Reference** check box.

7. Click **Save Security Settings**.

8. Click **Confirm**, and then click **Yes**.

The security settings of the controller are saved.

Disable security key

Prerequisites

- All secured virtual disks and non-RAID disks must be deleted or removed to disable security.
- All secured disks must be cryptographically erased.
- Any Auto Secure non-RAID options must be disabled.

About this task

The **Disable Security Key** feature is available only if a security key present on the controller.

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Enter the PERC 12 HII Configuration Utility](#).

2. Click **Main Menu > Controller Management > Advanced Controller Management > Disable Security**.
You are prompted to confirm whether you want to continue.

3. Click **Confirm**.

4. Click **Yes**.

The security key is disabled successfully.

NOTE: All secured disks must be erased or removed to disable security.

Create a secured virtual disk

About this task

To create a secured virtual disk, you must first create security key for the controller.

NOTE: Do not mix the following when creating a secured VD:

- SAS and SATA drives
- HDDs and SSDs
- NVMe drives with HDDs and SSDs
- TCG Enterprise and TCG Opal SED protocols

NOTE: To disable the security features, you must disable the Auto Secure Configuration settings.

After creating the security key, do the following:

Steps

1. Enter the Dell PERC 12 Configuration Utility. See [Enter the PERC 12 HII Configuration Utility](#).
2. Click **Main Menu > Configuration Management > Create Virtual Disk**.
3. Select the **Secure Virtual Disk** option.
4. Click **Create Virtual Disk**.
The secure virtual disk is created successfully.

Secure a non-RAID disk

In HII, secure a non-RAID disk by using the security key of the controller.

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Enter the PERC 12 HII Configuration Utility](#).
2. Click **Main Menu > Device Management > Logical Enclosure <enclosure Number>**. An enclosure number can be 0, 1, 2, 3, 4, and so on.
All the physical disks associated with the selected enclosure are listed.
3. Select a non-RAID disk.
4. From the **Operations** drop-down menu, select **Secure Non-RAID Disk**.

Secure a pre-existing virtual disk

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Enter the PERC 12 HII Configuration Utility](#).
2. Click **Main Menu > Virtual Disk Management**.
The list of virtual disks is displayed.
3. Select a virtual disk.
4. From the **Operations** drop-down menu, select **Secure Virtual Disk**.

NOTE: The virtual disks can be secured only when the virtual disks are in Optimal state.

Import a secured non-RAID disk

If you are inserting a non-RAID disk into a system that has a controller key different from the security key on the drive, the security key from the system in which it was initially secured must be provided in HII.

Prerequisites

 **NOTE:** The controller must have an existing security key before importing a secured non-RAID disk.

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Enter the PERC 12 HII Configuration Utility](#).
2. Click **Main Menu > Configuration Management > Manage Foreign Configurations**.
3. Select **Enter Security Key For Locked Drives** and enter the security key if you are importing drives with a different security key.
The drive's key will be changed to the controller key.
4. If required, convert the drive to a non-RAID drive. See [Convert a physical disk to a non-RAID disk](#).

Import a secured virtual disk

Prerequisites

 **NOTE:** The controller must have an existing security key before importing secured foreign virtual disk.

Steps

1. Enter the **Dell PERC 12 Configuration Utility**. See [Enter the PERC 12 HII Configuration Utility](#).
2. Click **Main Menu > Configuration Management > Manage Foreign Configurations**.
3. Click **Import Foreign Configuration**.
4. To import a VD by using a different security key, enter the new security key.
5. When prompted to conform if you want to perform the operation, click **Confirm**.
6. Click **Yes**.

The foreign configuration is imported successfully.

 **NOTE:** By using PERC 12, you cannot import foreign configuration data of previous PERC versions.

Dell Technologies OpenManage Secure Enterprise Key Manager (SEKM)

This feature allows the PERC to receive a security key from a remote server instead of saving the key on a local controller. This protects data on secured disks under the PERC if the disks or the entire system is stolen. For more information about configuring OpenManage Secure Enterprise Key Manager (SEKM) and Sockets Layer (SSL) or Transport Layer Security (TLS) related configuration, see the iDRAC User's Guide available on the support site. See the Enable OpenManage Secure Enterprise Key Manager (SEKM) on PowerEdge Servers technical white paper available on the Dell support site.

 **NOTE:** When replacing a controller enabled with enterprise key management, you must re-enable enterprise key management on the controller from iDRAC.

 **NOTE:** If key exchange fails during boot, view and correct any connection issues with the key server that is identified in the iDRAC Lifecycle log, and then cold reboot the server.

 **NOTE:** VDs will no longer be automatically rekeyed when they are unlocked. They will be rekeyed only when they are imported. Non-RAID and unconfigured drives will still be automatically rekeyed.

 **NOTE:** The Relock property is available in NVMe drives for both native and foreign locked drives and SAS/SATA for foreign locked drives.

NOTE: Support for reconfiguring NVMe drives is allowed for PDs and other drives that have to be erased to change the format.

Supported controllers for OpenManage Secure Enterprise Key Manager (SEKM)

Enterprise key manager mode is supported on the PERC H965i Adapter, PERC H965i Front, PERC H965i Front DC-MHS PERC H965i MX, and PERC H965e Adapter controllers, and allows the creation of secured virtual disks and non-RAID disks. For more information about supported platforms, see the iDRAC User's Guide available on the support site.

NOTE: The SEKM mode is not supported on PERC H965i MX.

Manage the Server Enterprise Key Manager (SEKM) feature

iDRAC manages the SEKM features. For instructions on enabling enterprise key manager mode, see the SEKM section in the relevant version of the iDRAC User's Guide available on the support site.

NOTE: When the SEKM mode is enabled, the controller waits up to two minutes for iDRAC to send keys, and then the PERC continues to boot.

NOTE: iDRAC performs rotation of keys. Any attempt to rekey the controller through a different management application is not supported.

Disable the Secure Enterprise Key Manager (SEKM) feature

The SEKM mode can be disabled only by using the iDRAC interfaces such as GUI and RACADM.

For information about disabling SEKM, see the "SEKM Functionalities" section in the relevant iDRAC User's Guide available on the support site. Also see the Enable OpenManage Secure Enterprise Key Manager (SEKM) on Dell PowerEdge Servers technical white paper available on the support site.

Manage virtual disks in the SEKM mode

The VDs in SEKM mode are managed in the same way as in local key manager mode. SED capable VDs can be secured during or after creation. See [Create a secured virtual disk](#).

Manage non-RAID disks in the SEKM mode

Non-RAID disks are managed in the same way in enterprise key manager mode as in local key manager mode. SED capable non-RAID disks can be secured after creation. See [Secure a non-RAID disk](#).

Transition of drives from Local Key Management to Secure Enterprise Key Management

PERC enables transition from Local Key Management (LKM) mode to Secure Enterprise Key Manager (SEKM) mode without disabling LKM security first. For instructions on transitioning from LKM mode to SEKM mode, see the iDRAC User's Guide of the relevant iDRAC version available on the support site.

You cannot transition from LKM to SEKM when:

- The snapdump is present on PERC.
- The Sanitize operation on a physical disk is in progress.
- The LKM key does not match with the current key of PERC.

Troubleshooting issues in PERC 12 series cards

To get help for resolving issues in your Dell Technologies PowerEdge RAID Controller 12 series cards, you can contact your Dell Technical Service representative.

Topics:

- Virtual drive Blink and Unblink operations are unavailable in iDRAC UI and HII
- Samsung PM893 and Samsung PM897 SATA SSDs are not supported on PERC12 firmware versions earlier than 8.8.0.0.18-26
- The SMART data of SATA SSDs is unavailable from iDRAC
- The default HDD cache policy is not updated to support the read-ahead operation
- Firmware downgrade operation may fail if the firmware version downgraded to does not have the supported feature set
- Unable to downgrade the firmware version of PERC12 to an older revision
- When secured drives are hot-removed or hot-inserted in the server, iDRAC doesn't generate the PD Unlock event
- During a Key Exchange Failure scenario in SEKM mode, native locked drives are Identified as "Locked with Foreign Key" in all interfaces
- After importing the foreign configuration of the associated global hot-spare drive on a PERC12 series controller, the drive is still reported as both Online and global hot-spare
- LED Operations using ESXCLI Commands for H965 and HBA465 Series controller is not available using VMware ESXi OS
- Backplane slot count may be reported as 32 if there is an error with the backplane during discovery or runtime
- Single virtual disk performance or latency in hypervisor configurations
- Unable to discover or detect a PERC card
- Unable to install the Microsoft Windows operating system
- Only UEFI is supported and not BIOS
- A fault firmware state is detected
- Foreign configuration not found in HII
- Degraded state of virtual disks
- Boot-time errors
- Event log errors
- Application issues
- Memory errors
- Preserved Cache state
- Security key errors
- General issues
- Physical disk issues
- SMART errors
- Replace member errors
- Linux operating system errors
- Drive indicator codes
- The AutoSecure feature is not supported
- HII error messages

Virtual drive Blink and Unblink operations are unavailable in iDRAC UI and HII

Error Message: VD Blink and Unblink operations are greyed out on the iDRAC user interface.

Corrective Action: Known issue. No corrective action required. Use the locate command: `perccli2 /c0/e32/s0 start locate`. For more information, see the "Locate a drive" section in the PERC12 CLI Reference Guide available on the support site.

Samsung PM893 and Samsung PM897 SATA SSDs are not supported on PERC12 firmware versions earlier than 8.8.0.0.18-26

Error Message: Samsung PM893 and Samsung PM897 SATA SSDs are not detected by the controller.

Corrective Action: Update to PERC12 8.8.0.0.18-26 and later firmware versions.

The SMART data of SATA SSDs is unavailable from iDRAC

Error Message: The SMART data of SATA SSDs is unavailable from iDRAC when it is connected to a PERC12 or HBA465 series controller.

Corrective Action: Known issue. Use the PERCCLI2 to view the SMART data about the SATA SSDs.

The default HDD cache policy is not updated to support the read-ahead operation

Error Message: The default HDD cache policy is not updated to support the read-ahead operation.

Corrective Action: Any newly created HDD VDs will not have read-ahead enabled by default. You must manually enable using the PERC CLI.

Firmware downgrade operation may fail if the firmware version downgraded to does not have the supported feature set

Error Message: A validation error occurred during the firmware update indicating the FEATURE COMPATIBILITY CHECK failure event. For example, inability to downgrade the firmware version of PERC12 H965i or H965e controller because the read-ahead feature is enabled on the VD.

Corrective Action: Expected behavior if the firmware behavior is 8.8.0.0.18-26 and later versions. If enabled, disable the read-ahead feature on the VD before downgrading the firmware version.

 **NOTE:** Dell recommends that you must always use the latest firmware version of the controller available on the support site.

Unable to downgrade the firmware version of PERC12 to an older revision

Error Message: Unable to downgrade the firmware version of PERC12 to an older revision.

- Downgrade to 8.4.0.0.18-xx or older firmware is blocked if:
 - the Read Ahead feature is enabled on a VD.
 - secured disks are present.
- Downgrade to 8.0.0.0.18-xx or earlier version is blocked if:
 - An NVMe switch is connected to the controller.
 - VD count is greater than the supported limit .
 - Online Capacity Expansion (OCE) is enabled on the VD.
 - The PD Erase feature is enabled.
- Any upgrade or downgrade operation is blocked if:
 - Preserved cache is present.

Probable cause: Expected behavior.

Corrective Action Check the controller event log for the validation check failure. Disable any new features or wait for any OCE or erase operation to finish on the controller, VDs, or PDs, and then retry the downgrade operation. If preserved cache is present then insert missing disks or delete the preserved cache and retry the downgrade operation.

When secured drives are hot-removed or hot-inserted in the server, iDRAC doesn't generate the PD Unlock event

Error Message: When secured drives are hot-removed or hot-inserted in the server, iDRAC doesn't generate the PD Unlock event.

Probable Cause: The PD Unlock event is happening before the PD Insert event.

Corrective Action: Known issue. No response action is required. There is no functional impact. The secure drives will be already in unlock state.

During a Key Exchange Failure scenario in SEKM mode, native locked drives are Identified as "Locked with Foreign Key" in all interfaces

Error Message: When key exchange fails in the SEKM mode, the native locked drives are identified as "Locked with foreign key".

Probable Cause: Could be one of the following: 1) An issue in connecting to the SEKM server. 2) A key is deleted from the KMS.

Corrective Action: Known issue. Check the SEKM server connectivity to iDRAC.

After importing the foreign configuration of the associated global hot-spare drive on a PERC12 series

controller, the drive is still reported as both Online and global hot-spare

- Error Message:** After importing the foreign configuration of the associated global hot-spare drive on a PERC12 or HBA465 series controller, the drive is still reported as both Online and global hot-spare
- Probable Cause:** The drive type of a commissioned Global Hot Spare was incorrectly set as Global Hot Spare during the foreign import operation.
- Corrective Action:** Reset the controller by running the `perccli2 /c0 reset` command at the PERC12 CLI.

LED Operations using ESXCLI Commands for H965 and HBA465 Series controller is not available using VMware ESXi OS

- Error Message:** The `- esxcli storage core device physical get -d naa.58ce38ee2182900d` command used to locate the drive by blinking or unblinking the LED when using an ESXi7 or ESXi8 operating system.
- Probable Cause:** The LSU Plugin does not have SL8 embedded to support the blink and unblink operations.
- Corrective Action:** Known Issue with a VMware ESXi Operating System. Use the `locate` command: `perccli2 /c0/e32/s0 start locate`. For more information, see the "Locate a drive" section in the PERC12 CLI Reference Guide available on the support site.

Backplane slot count may be reported as 32 if there is an error with the backplane during discovery or runtime

- Error Message:** Backplane is indicated as faulty and slot-count is shown as 32 even though the configuration supports less than 32.
- Corrective Action:** Known issue. No functional impact. System will require a reset to recover from this behavior. For more information about backplanes, see the Installation and Service Manual (ISM) of the respective server model available on the support site.

Single virtual disk performance or latency in hypervisor configurations

Multi-initiator or hypervisor configurations running multiple I/O workloads to a single RAID array may experience degraded performance or latency. This is caused by upper layers sending separate I/O workloads for each virtual machine to the storage subsystem which ends up being a random I/O workload to the underlying RAID array. For I/O workload configurations that require lower latency restrictions and higher I/O performance it may be beneficial to run fewer I/O workloads to individual RAID arrays or to use separate RAID arrays and physical disks for each I/O workload. Other considerations are making sure write-back cache is enabled for rotational disks or using solid state drives (SSDs) to improve random I/O workload performance.

Performance degradation may also be observed when background operations such as initialization, consistency check, or rebuilds are running on the virtual disk. See your hypervisor storage best practices or performance best practices guides for additional configuration support.

Unable to discover or detect a PERC card

- Error Message:** A discovery error has occurred, please power cycle the system and all the enclosures attached to this system.
- Probable Cause:** This message indicates that disk topology discovery did not complete within 120 seconds. The cables from the PERC controller to the backplane might be improperly connected.
- Corrective Action:** Check the cable connections and fix any problems. Restart the system.

Unable to install the Microsoft Windows operating system

Ensure that you perform the following step before installing Windows on 4 KB sector drives:

1. Read and understand the updates to the version of Windows that you have installed. You can find this information in the Microsoft help. For more information, see [Microsoft support policy for 4 K sector hard drives in Windows](#).

Only UEFI is supported and not BIOS

- Issue** Only UEFI boot is supported by HBA465 and PERC12 controllers. BIOS is not supported.
- Corrective Action:** Expected behavior. No action is required.

A fault firmware state is detected

- Error Message:** Firmware is in Fault State. Controller: Broadcom MPI3 I/O Controller (Fault). The controller is in a Fault state: Faultcode <FaultCode>; Additional code: <Code> : <Code> : <Code>
- Corrective Action:** Contact your Technical Support team.

Foreign configuration not found in HII

- Error Message:** The foreign configuration message is present during POST but no foreign configurations are present in the foreign view page in HII configuration utility. All virtual disks are in an optimal state.
- Corrective Action:** Ensure all your PDs are present and all VDs are in optimal state. Clear the foreign configuration using **HII configuration utility** or **Dell OpenManage Server Administrator Storage Management**.
-  **CAUTION: The physical disk goes to Ready state when you clear the foreign configuration.**

If you insert a physical disk that was previously a member of a virtual disk in the system, and that disk's previous location has been taken by a replacement disk through a rebuild, you must manually remove the foreign configuration flag of the newly inserted disk.

Degraded state of virtual disks

A redundant virtual disk is in a degraded state when one or more physical disks have failed or are inaccessible. For example, if a RAID 1 virtual disk consists of two physical disks and one of them fails or becomes inaccessible, the virtual disk becomes degraded.

To recover a virtual disk from a degraded state, you must replace the failed physical disk and rebuild it. Once the rebuilding process is complete, the virtual disk state changes from degraded to optimal.

Boot-time errors

The following table lists error messages, probable causes, and recommended response action to resolve the issue.

Table 20. Boot-time issues and corrective actions

Description	Error Message	Probable Cause	Recommended Corrective Action
Removed or Missing Device	Some configured disks have been removed from the system or are no longer accessible. Check the cables and ensure all drives are installed in the slots.	Some configured drives are removed. Else, they may not be accessible anymore. Or, the cables from PERC to the backplane may be improperly connected.	Check the cables and ensure all drives are installed in the slots. If there are no cable related issues, restart the server.
Removed or Missing Device	The following VD's are missing <VD Names>. If you proceed or continue to start the Configuration Utility, these VD's will be removed from the configuration. To use VD's at a later time, you must import them.	Because some configured drives are removed. Else, they may not be accessible anymore. Or, the cables from PERC to the backplane may be improperly connected.	Check the cables and ensure all drives are installed in the slots. If there are no cable related issues, restart the server.
Removed or Missing Device	The following VD's are missing <VD Names> complete spans. If you proceed or continue to start the Configuration Utility, these VD's will be removed from the configuration. To use VD's at a later time, you must import them.	Because some configured drives are removed. Else, they may not be accessible anymore. Or, the cables from PERC to the backplane may be improperly connected.	Check the cables and ensure all drives are installed in the slots. If there are no cable related issues, restart the server.
Removed or Missing Device	All drives associated with the previous configuration are no longer associated with the configuration.	Because some configured drives are removed. Else, they may not be accessible anymore. Or, the cables from PERC to the backplane may be improperly connected.	Check the cables and ensure all drives are installed in the slots. If there are no cable related issues, restart the server.
Offline VD	The following VD's have missing PD's: <VD_Names>. If you continue or start the Configuration Utility, these VD's will be identified as Offline, and may become inaccessible.	Some configured drives are either removed or have stopped functioning, resulting in the VD's to be identified as Offline.	Check the cables and ensure all drives are installed in the slots. If there are no cable related issues, restart the server.
VD in write-through mode	The write-back VD's are temporarily running in write-through mode.	The controller energy pack may have low voltage or may not be functional. It can also be because the energy pack or super capacitor being charged is either missing or not in good health state.	Ensure that the cable connection to the energy pack cable is correct. Ensure the battery (Energy Pack) health status is good. If you are using an energy pack, please allow the energy pack to charge for 24 hours before evaluating the energy pack for replacement. You can evaluate the health of the energy pack or super capacitor by using the appropriate utility within the operating system or within POST.
Preserved Cache	Some VD's with preserved cache are either offline or not installed.	VD's had dirty cache during a server boot operation.	Insert the missing drives and restart the server. Check the cables and ensure all drives are installed in the slots.

Table 20. Boot-time issues and corrective actions (continued)

Description	Error Message	Probable Cause	Recommended Corrective Action
The Foreign Configuration Import operations did not import any drives.	The Foreign Configuration Import operations did not import any drives.	Either the Foreign Configuration is incompatible with this controller or the Foreign Configuration is incomplete.	Make sure that either the drives necessary for this configuration are installed or remove the incompatible configuration.
Cache Discarded	Some memory or energy pack issues were detected. An unexpected power loss has occurred. The adapter has been recovered, but the controller cache was lost.	The battery (Energy Pack) may be not in good state or has low voltage during the AC cycle of a server that has dirty cache.	Contact the Dell Technical support team.
Cache Discarded	The controller cache was discarded because of an unexpected power-off or reboot operation during a write operation, but the adapter has recovered. Or, some memory or energy pack issues were detected. An unexpected power loss has occurred. The adapter has been recovered, but the controller cache was lost.	The battery (Energy Pack) may be not in good state or has low voltage during the AC cycle of a server that has dirty cache. The DDR4 device on the server may be moving to a nonfunctional state.	Contact the Dell Technical support team.
Cache Discarded	An issue is observed in restoring the offloaded cache. The data cache is lost.	The battery (Energy Pack) may be not in good state or has low voltage during the AC cycle of a server that has dirty cache.	Contact the Dell Technical support team.
Cache Discarded	After resetting, the Onboard Controller Memory (OCM) data restore operation is unsuccessful.	The onboard controller memory may either have an issue or moved to a bad state.	Contact the Dell Technical support team.
Cache data lost because of Bad or Low Voltage Energy Pack	Cache Offload was not sustainable. This could be caused by bad or low voltage energy pack or battery. Write journals could have been possibly lost with parity-based arrays (RAID 5, 6, 50, or 60).	Energy pack is not charged fully.	Wait for the energy pack to charge back before enabling the WB/AWB cache. If issue persists, replace the energy pack.
Multi-Bit ECC errors detected	Multi-bit ECC errors were detected on the RAID controller.	The DDR4 device on the server may be moving to a nonfunctional state.	Controller may need to be replaced. Contact Dell Technical support team.
Single-Bit ECC errors detected	Single-bit ECC errors were detected during the previous boot operation of the RAID controller.	The DDR4 device on the server may be moving to a nonfunctional state.	Controller may need to be replaced. Contact Dell Technical support team.
Too many bad block to support data backup	The Non-Volatile Cache storage capacity is too low to support data backup. The write-back VDs will be converted to write-through. Or, Non-Volatile Cache has gone bad.	The Open NAND Flash Interface (ONFI) flash used for DDR backup is not functioning.	Contact the Dell Technical support team.
Non-Volatile Cache Degraded	The data backup capacity of the Non-Volatile Cache device is degraded.	The Open NAND Flash Interface (ONFI) flash used for DDR backup is not functioning.	Consider replacing the device or contact the Dell Technical support team.
Factory Setting Corrupt	The non-volatile data validation operation is unsuccessful.	The factory settings of the controller may have been corrupted.	Upgrade the firmware by using the correct NV Data. Or, contact the Technical Support team.
Factory Setting Corrupt	Unable to read the MPB file of the personality, OEM ID, or Profile ID. The controller is switching to Safe Mode.	The factory settings of the controller may have been corrupted.	Contact the Dell Technical support team.

Table 20. Boot-time issues and corrective actions (continued)

Description	Error Message	Probable Cause	Recommended Corrective Action
Factory Setting Corrupt	The NVRAM layout is either corrupted or has a mismatch, and therefore, is reinitialized.	The factory settings of the controller may have been corrupted, and therefore, reinitialized.	Upgrade the firmware by using the correct NV Data. Or, contact the Technical Support team.
Enclosure/ Backplane Count exceeded	The number of enclosures or backplanes connected to the connector <connectorName> has exceeded the maximum number.	Backplanes or enclosures connected to the connector are more than the maximum allowed.	Power off the server and remove the backplanes or enclosures to ensure that the number is less than the maximum value.
Physical Disk Count exceeded	The number of PDs connected are more than the supported value of <number> drives.	The installed PDs are more than the maximum allowed.	Power off the server and remove the PDs to ensure that the number is less than the maximum value.
Topology Error	An invalid SAS topology is detected in <name>.	Either a non-functional or corrupted PD is inserted in the server.	Remove any recently inserted PDs from the server.
Discovery Error	Unable to discover the controller: <%s>.	The controller could not be discovered within 120 seconds. Cables from the PERC to the backplane might be improperly connected.	Check the cable connection between PERC and server. Power cycle the server and all attached enclosures.
Unable to communicate with Key management server	Unable to communicate with the SEKM Server. If you continue, there will be a drive security key error and all the secured configurations will be marked as foreign. Please check the connection with the SEKM server, reboot the machine to retry switching to EKM.	Connection Information between iDRAC and the Key Management Server (KMS) may have changed. Or, An iDRAC controller discovery issue may have occurred preventing the key exchange within the timeout period.	See the iDRAC User's Guide of the relevant version on the support site, ensure that the KMS communication is successful, and then restart the server.
Safe Mode entered	The controller booted in the safe mode.	An internal issue forced the controller to boot in the safe mode.	View the server screen to get information about why the controller booted to a safe mode, and then take the recommended corrective action.
Safe Mode Exited	The controller has exited the safe mode.	Not applicable.	No response action is required.

Event log errors

The following table lists error messages, probable causes, and recommended response action to resolve the Event Log errors.

 **NOTE:** Text used in the following messages may vary based on the type of management application.

Table 21. Run-time issues and corrective actions

Description	Error Message	Probable Cause	Recommended Corrective Action
Correctable Error During BGI	A medium error was corrected on the following virtual drive during background initialization: <VD_Name> at <variable Name>, <PD_Name>, LUN <Name>, Count <Number>.	The drive had a correctable medium error and the data is recovered. The drive may be getting corrupted or facing a data-rot.	Verify the SMART status of the drive, and if necessary, replace the drive.

Table 21. Run-time issues and corrective actions (continued)

Description	Error Message	Probable Cause	Recommended Corrective Action
Uncorrectable Error During BGI	The Background Initialization operation is completed in the virtual drive <VD_Name> with uncorrectable errors.	One or more uncorrectable errors are detected during the BGI operation. Data may be lost.	Verify the SMART status of the drive and ensure that the data is backed up.
Uncorrectable Error During BGI	The Background Initialization operation is completed in the virtual drive <VD_Name> with uncorrectable errors.	An uncorrectable error was detected during the BGI operation. Data may be lost.	Verify the SMART status of the drive and ensure that the data is backed up.
BGI Failed	Unable to complete the Background Initialization operation on the virtual drive <VD_Name>.	A physical drive may have failed during the BGI operation which made the virtual drive appear as offline.	Check the health status of virtual drives and physical drives, and then replace corrupted drives if necessary.
Error Corrected during consistency check	The Consistency Check operation corrected error on the drive: <VD_Name> at <variable Name>, <PD_Name>, LUN <Name>, Count <Number>.	The drive had a medium error and corrected during Consistency Check. Data may be lost.	Verify the SMART status of the drive and ensure that the data is backed up.
Consistency check completed with corrected errors	The Consistency Check operation was completed on the drive: <VD_Name> Corrections <correctionData>.	The drive had a correctable medium error and the data is recovered. The drive may be getting corrupted or facing a data-rot.	Verify the SMART status of the drive, and if necessary, replace the drive.
Uncorrectable media errors detected during consistency check	The Consistency Check operation detected uncorrectable multiple medium errors on the drive: <VD_Name> at <variable Name>, <PD_Name>, LUN <Name>, Count <Number>.	The drive had a medium error and corrected during Consistency Check. Data may be lost.	Check the health status of virtual drives and PDs, and then replace corrupted drives if necessary.
Unable to complete consistency check	Unable to complete the Consistency Check operation on the virtual drive <VD_Name>.	Consistency Check has failed on virtual drive which could be because of media errors.	Check the health status of the virtual drives and PDs. Replace any non-functional drives.
Consistency check completed with uncorrectable errors	The Consistency Check (CC) operation is completed with uncorrectable errors on the virtual drive <VD_Name>.	Medium Errors were not corrected after CC. If virtual drive is not redundant (or is degraded), then data cannot be regenerated and the medium error cannot be resolved. Data may be lost.	Check the health status of virtual drives and physical drives, and then replace corrupted drives if necessary.
Inconsistent parity data detected during consistency check	The Consistency Check operation detected inconsistency parity on the virtual drive <VD_Name>, at strip <stripName>.	The Consistency Check operation detected an on the virtual drive. The inconsistency is corrected.	No response action is required.
Consistency check logging disabled due to too many inconsistencies	Inconsistency data on the virtual drive <VD_Name> cannot be logged in because too many inconsistencies are detected on the virtual drive, and the feature is disabled during the Consistency Check operation.	The Consistency Check operation detected multiple inconsistencies and disabled data logging.	Check the health status of virtual drives and physical drives, and then replace corrupted drives if necessary.
Virtual drive Initialization failed	Unable to initialize the virtual drive <VD_Name>.	A disk could have failed during initialization causing the virtual drive to go offline.	Check the health status of virtual drives and physical drives, and then replace corrupted drives if necessary.

Table 21. Run-time issues and corrective actions (continued)

Description	Error Message	Probable Cause	Recommended Corrective Action
Physical drive erase errors	Data on the following physical drive cannot be cleared: <PD Name>, <Path Name>, <ErrorDescription>.	The drive could have either failed or returned an error during the Clear operation.	Check the health and SMART status of physical drives, and then replace corrupted drives if necessary.
Physical drive error	An error is detected on the following physical drive: <PD Name>, <Path Name>, <ErrorDescription>.	An error is detected on the physical drive.	Check the health and SMART status of physical drives, and then replace corrupted drives if necessary.
Physical drive not supported	The physical drive <PD Name> is either not supported by the controller or is in an unsupported format.	The physical drive is either not supported by the controller or is in an unsupported format.	Replace the drive or reformat by using a supported file format.
Physical drive not certified warning	The physical drive <PD Name> is not certified.	The physical drive is not manufactured as per Dell technical specifications and cannot guarantee that the physical drive will be fully compliant with Dell standards and functionalities.	Use a physical drive that complies to Dell standards and retry the operation.
Media error corrected during patrol-read	The patrol-read operation corrected a medium error on the physical drive <PD Name>, LUN <LUN Name>.	The drive had a medium error and corrected during Consistency Check. Data may be lost.	Verify the SMART status of the drive and ensure that the data is backed up.
Uncorrectable media error detected during patrol-read	The patrol-read operation detected an uncorrectable medium error on the physical drive <PD Name>, LUN <LUN Name>.	An uncorrectable error was detected during the patrol-read operation. Data may be lost.	Check the health status of virtual drives and physical drives, and then replace corrupted drives if necessary.
Physical drive Predictive Failure	A predictive failure is detected in the physical drive: <PD Name>.	The SMART data is indicating that the drive may fail soon.	Replace the drive and ensure the rebuild or copy-back operation is successful.
Bad block punctured on the physical disk	A bad block is being punctured on the physical drive: <PD Name>, LUN <LUN Name>.	LBA on physical drive was punctured. Data may be lost.	Check the health status of virtual drives and physical drives, and then replace corrupted drives if necessary.
Unable to rebuild array because of an error on the source disk	Unable to rebuild the physical drive <PD Name> because of an error in the source drive.	Source drive failed because of which the Rebuild operation cannot proceed. Data may be lost	Replace any bad drives and recreate virtual drive from backup.
Unable to rebuild array because of an error on the target disk	Unable to rebuild the physical drive <PD Name> because of an error in the target drive.	Target drive failed due to which the Rebuild operation cannot be continued.	Replace any bad drives and restart the Rebuild operation.
Unrecoverable error detected during rebuild	An unrecoverable medium error is detected during the Rebuild operation on the physical drive: <PD Name>, LUN <LUN Name>.	An uncorrectable error was detected during the rebuild. Data may be lost.	Check the health status of virtual drives and PDs, and then replace corrupted drives if necessary.
Media error corrected during operation	A medium error is corrected during the recovery operation on the physical drive: <PD Name>, LUN <LUN Name>.	An medium error was detected during the BGI operation. Data may be lost.	Verify the SMART status of the drive and ensure that the data is backed up.

Table 21. Run-time issues and corrective actions (continued)

Description	Error Message	Probable Cause	Recommended Corrective Action
Unrecoverable error detected during operation	An unrecoverable medium error is detected during the recovery operation on the physical drive: <PD Name>, LUN <LUN Name>.	Medium Errors were not corrected after recovery. If a virtual drive is not redundant (or is degraded), data cannot be regenerated and medium error cannot be resolved. Data may be lost.	Check the health status of virtual drives and physical drives, and then replace corrupted drives if necessary.
SCSI Sense Data	An unexpected Sense error is on the physical drive: <PD Name>, Patj <Path Name>, CDB <CDB Name>, Sense <Sense Name>.	Drive returned sense data for the SCSI CDB. This can include info during drive power-on, firmware update, or errors.	If the event is seen consistently in logs, contact the Technical Support team.
Physical drive is not accessible	Unable to access the PD: <PD Name>.	An error occurred during the discovery or initialization of the drive. The drive may be corrupted.	Remove or reinstall the drive. Check the health or SMART status of the physical drive. For NVMe drives, run the Recovery operation of the supported format.
The hot-spare is no longer useful	The dedicated hot-spare <PD Name> is no longer useful for all the arrays.	The dedicated hot-spare cannot cover all the arrays and this could be because of a change in the virtual drive size or other parameters.	Check if the assigned dedicated hot-spare physical drive is compatible with the virtual drive it covers. If not, assign a hot-spare physical drive which is compatible.
Hot-Spare no longer covers all arrays	The global hot-spare <PD Name> no longer covers all the arrays.	The assigned global hot-spare physical drive is not of the same type to cover all the virtual drives in the system.	Ensure to assign a global hot-spare physical drive which is of the same type of the other VD. Else, use dedicated hot-spare for virtual drives that are not covered.
Energy Pack is not installed	The Energy Pack is not installed.	Energy pack may be either missing, or have a disconnected or damaged cable.	Ensure that the Energy Pack is installed, connected through a cable. Else, replace the Energy Pack.
Energy Pack temperature is high	The Energy Pack temperature is high.	Energy pack temperature exceeded the operational limit.	Verify system environment is within operational limits. Check for fan failure, or air flow blockage. In the absence of any other temperature condition, replace energy pack.
Energy Pack temperature is low	The Energy Pack temperature is low.	Energy pack may self-discharge in storage, which should self-correct in operation within a few minutes. Virtual drives may transition from write-back to write-thru until condition is resolved.	Persistent condition may require energy pack or controller replacement. Contact support.
Energy Pack temperature is not functional	The Energy Pack is not functioning and cannot support data retention.	Energy pack has failed. Virtual drives will transition to the write-through mode until corrected.	Replace the Energy Pack.
Energy pack below charge threshold	The current capacity of the Energy Pack is less than the threshold value.	May result from energy pack learning cycle indicating it is approaching end of life. Virtual drives will transition to	Replace the Energy Pack.

Table 21. Run-time issues and corrective actions (continued)

Description	Error Message	Probable Cause	Recommended Corrective Action
		the write-through mode until corrected.	
Energy pack above charge threshold	The current capacity of the Energy Pack is more than the threshold value.	May result from energy pack learning cycle.	No response action is required.
Unable to communicate with backplane or enclosure	The controller is unable to communicate with the enclosure <enclosure Name>.	The cable may be loose or damaged.	Ensure that the cables are connected and not damaged. Restart the server if necessary.
Backplane or enclosure discovery error	Discovery error detected for enclosure <enclosure name> (receptacle <receptacle #>) - <error code>.	Error is probably caused by cables, backplane firmware, or slot connections.	Check all cables and drives and ensure they are properly seated. AC power-cycle the system. If the issue persists, contact your Technical Support team.
PHY/Slot bad	The enclosure with physical drive <PD Name> is not detected by controller in the slot <slotName>.	The physical drive is inserted in the slot but not detected by the controller. The physical drive may be corrupted.	Ensure that the cables are connected and not damaged. Restart the server if necessary.
The backplane or enclosure is unstable.	The enclosure <enclosure Name> is unstable.	The backplane or enclosure has detected an error.	Ensure that the cables are connected and not damaged. Upgrade the enclosure firmware and restart the server.
Hardware issues detected on backplane or enclosure	The enclosure <enclosure Name> is has hardware issue.	The backplane or enclosure has indicated that there is a hardware failure or issue.	Ensure that the cables are connected and not damaged. Upgrade the enclosure or backplane firmware and restart the server. If the issue persists, contact your Technical Support team.
Enclosure is not responding	The enclosure with physical drive <enclosure Name> is not responding.	The backplane is not responding to the controller.	Ensure that the cables are connected and not damaged. Upgrade the enclosure or backplane firmware and restart the server. If the issue persists, contact your Technical Support team.
Virtual drives transitioned to write-through because of issues in the energy pack hardware	The Energy pack charger has been disabled because of hardware issues. Currently changing the write-back virtual drives to write-through. However, the forced write-back virtual drives are not affected.	The energy pack is not functional or defective. Therefore, the virtual drive cache policy is configured for forced write-back.	Verify the energy pack connection, check for cable damage, replace the energy pack. If the issue persists, contact the Technical Support team.
Issues detected with energy pack or charger	Issues are detected in either the Energy Pack or Energy Pack Charger.	The State of Health (SoH) of the Energy Pack is not good. Energy pack is bad, defective, missing, or there is a fault in the energy pack charging system. Virtual drives in the write-back mode will transition to write-through.	Verify the energy pack connection, check for cable damage, replace the energy pack. If the issue persists, contact the Technical Support team.
Disk missing	The following physical drive is missing: <PD Name>.	A configured physical drive is not detected after resetting or restarting the controller.	If this issue is unexpected, ensure that the cables are properly connected and physical

Table 21. Run-time issues and corrective actions (continued)

Description	Error Message	Probable Cause	Recommended Corrective Action
			drives are seated, and then restart the server.
Virtual disk offline because of missing disks	The following virtual drives are not detected, and therefore, the physical drives will go offline: <VD Names>.	One or more of the physical drives is bad and has gone offline.	Check the health status of the physical disks and replace any bad drives.
Virtual drive is partially degraded	The virtual drive <VD Name> is partially degraded.	One or more of the physical drives is bad and has gone offline.	Check the health status of the physical disks and replace any bad drives.
Virtual drive is degraded	The VD <VD Name> is now fully degraded.	One or more of the physical drives is bad and has gone offline.	Check the health status of the physical disks and replace any bad drives.
Virtual drive is offline.	The virtual drive <VD Name> is now offline.	One or more of the physical drives is bad and has gone offline.	Check the health status of the physical disks and replace any bad drives.
Physical drive command timeout	The command timed out waiting for response: <PD Name>, Patj <Path Name>, CDB <CDB Name>, Sense <Sense Name>.	The drive Command timed out waiting for response. This may occur due to an error or in cases where the device is reset such as a firmware update. Command may have been retried.	Check the health or SMART status of the drive and replace any failing drives. Try restarting the server or controller. If the issue persists, contact the Technical Support team.
Disk Reset	The physical drive <PD Name>, Path <pathname> is reset.	Drive was reset. This can happen in an attempt to recover from an error during device discovery.	If the issue is persistently seen in logs, contact the Technical Support team.
Bad Block table for virtual drive 80% full	The bad block table on the virtual drive <VD Name> is 80% full.	Multiple uncorrectable medium errors were found on the drive and added to LDBBM table.	Check the health status of the physical disks and replace any bad drives.
Bad block logging for virtual drive disabled because of too many bad blocks	Unable to log the block <blockName> on the physical drive <PD Name>, LUN <LUN Name> at <Count> because of a block table on the virtual drive <VD Name>.	Multiple uncorrectable medium errors were found on the drive and added to LDBBM table.	Check the health status of the physical disks and replace any bad drives.
Uncorrectable media error is detected on virtual drive	An uncorrectable medium error was logged for virtual drive <VD Name>: physical drive <PD Name>, LUN <LUN Name> at <Count> because of a block table on the virtual drive <VD Name>.	Uncorrectable multiple medium errors were found on the disk. Data may be lost.	Check the health status of the physical disks and replace any bad drives.
Media error corrected on virtual drive	A medium error was corrected on the virtual drive <VD Name> at <name>.	Medium Error was found on the disk and corrected. Data may be lost.	Verify the SMART status of the drive and ensure that the data is backed up.
Bad block table for virtual disk is 100% full	The bad block table on the virtual drive <VD Name> is 100% full.	Multiple uncorrectable medium errors were found on the drive and added to LDBBM table.	Check the health status of the physical disks and replace any bad drives.
Preserved Cache Discarded	The controller cache is discarded for either the deleted, missing, or offline virtual drive <VD Name>.	A state of a virtual drive has changed because drives are missing, failed, or deleted. Either the user has discarded the controller cache or the virtual drive was deleted.	Check the state of the virtual drive and all associated physical drives. Replace any non-functional drives.

Table 21. Run-time issues and corrective actions (continued)

Description	Error Message	Probable Cause	Recommended Corrective Action
Physical drives update timeout	The Microcode Update operation timed out on the physical drives <PD Name>.	The drive firmware update operation did not complete within the specified time limit. The firmware may still have updated successfully.	Verify if the drive firmware is fully updated. Else, retry the operation. If the firmware cannot be upgraded, then contact your Technical Support team.
Security key entered is invalid	Unable to unlock the drive because an invalid security key is entered.	An incorrect security key ID is used to unlock the drive or controller.	Enter a valid security key ID and retry the operation. If multiple drives use the same key ID and different security keys then remove the extra drives and separately import the drives.
Unable to unlock drives because of invalid security key	Unable to unlock the drive because the security key provided by Escrow is invalid.	The security key ID provided by Escrow is invalid.	If multiple drives use the same key ID and different security keys then remove the extra drives and import the disks separately.
A Drive Security Error is detected	Security subsystem issues are detected on the physical drive <PD Name>.	There is an issue when managing the security of the drive.	Remove, reinstall the drive, and then retry the operation. If the issue persists, contact your Technical Support team.
Controller cache preserved for VD	The controller cache data was preserved for the missing of offline virtual drive: <VD Name>.	The state of a virtual drive has changed because either drives are missing or failed.	Check the state of the virtual drive and all associated physical drives. Replace any non-functional drives.
Locked drives are not accessible because of a key exchange error	Unable to access the secured configuration because an incorrect security key is entered.	Either an incorrect security key is entered or a key is not entered at all. Locked drives will remain locked and cannot be accessed.	Using iDRAC, verify the SEKM status of the server. Resolve any network communication issues with the KMS server. To retry exchanging the security key, restart the server.
Invalid security key or key id entered	An invalid security key of the drive is entered.	The security key or key ID of the drive does not match the drive information.	Ensure that correct security key and key ID are entered to unlock the drive.
Unable to communicate with the external key manager	Unable to communicate with the external key manager.	Either an incorrect security key is entered or a key is not entered at all. Locked drives will remain locked and cannot be accessed.	Using iDRAC, verify the SEKM status of the server. Resolve any network communication issues with the KMS server. To retry exchanging the security key, restart the server.
Physical drive Erase Error	Unable to erase data on the physical drive: <PD Name>, Path <pathname>, Error Data <error Info>.	The drive may have an internal error or failed during the Erase operation.	Retry the operation. If the issue persists, contact your Technical Support team. If the data must be erased, then follow the DoD 5220.22-M Standard for Drive Erasure to physically destroy the storage device.
Energy pack charging disabled due to high temperature	The Energy Pack charging operation is abruptly stopped because the Energy Pack temperature is too high.	System environment is not within operational specifications, or energy pack is damaged.	Ensure that system environment is within operational limits. Check for fan failure, or air flow blockage around controller. In the absence of an environmental cause, disconnect energy pack and contact support.

Table 21. Run-time issues and corrective actions (continued)

Description	Error Message	Probable Cause	Recommended Corrective Action
Too many bad blocks to support cache backup	The non-volatile cache capacity is too less to support data backup. The write-back virtual drives will be converted to write-through.	Controller has run out of enough good cache buffer space to currently support caching.	Check the state of all virtual drives and all associated physical drives. Discard any preserved cache if necessary.
Non-Volatile cache has failed and backup cannot be supported	The non-volatile cache device failed and cannot support the data retention feature.	Controller cache support may not be functioning properly.	Contact your Technical Support team.
Controller temperature exceeded warning threshold	The controller temperature exceeded the threshold value. This may indicate inadequate server cooling. Currently switching to a low-performance mode.	The server fans may not be functioning properly.	Check the health of server fans and replace if necessary. To improve cooling, increase or offset the fan speed.
Controller temperature exceeded critical threshold	The controller has been shut down because the server temperature reached the threshold value. This indicates that the server cooling is inadequate.	The server cooling fans may not be functioning properly.	Increase or offset the fan speed to improve cooling.
Unable to secure virtual drive in future	Unable to secure the virtual drive <VD Name> in future because it is a non-SED.	One of the physical drives hosting the virtual drive is not SED-capable.	Replace the non-SED with an SED-capable one.
Discovery error	The drive cannot be discovered because of an SAS topology error: <Error Description>.	An error occurred during device discovery.	Remove all the newly added or non-functioning hardware. If the issue persists, contact your Technical Support team.
Energy pack unavailable or not functioning	The Energy Pack is either unavailable or not functioning. Incomplete write operations because of power loss may cause data integrity issues on the parity virtual drive: <VD Name>.	Energy pack is bad or defective, and the virtual drive cache policy is configured for forced write back.	Verify energy pack connection, check for cable damage, replace energy pack. Disable forced write back until resolved.
Safe Mode Error	The controller booted to the Safe mode because of critical issues.	Controller has an error during boot that requires user attention.	Correct the errors reported by the controller.
Firmware Update Failure	A validation issue occurred during the firmware update operation: <Issue Description>.	Failed to validate the incoming firmware image file.	Ensure that the firmware image used is correct and retry the operation. If the issue persists, contact your Technical Support team.
Firmware Update Failure	A programming issue occurred during the firmware update operation: <Issue Description>.	Programming of incoming firmware image to the flash failed.	Ensure that the firmware image used is correct and retry the operation. If the issue persists, contact your Technical Support team.
Firmware Update Failure	Unable to complete the Online activate operation while preparing the controller.	The new firmware may require offline activation.	Perform system reset to activate the firmware offline.
Security Key cleaned up	The Escrow key ID <key ID> is cleaned up.	The security key that was saved in the memory for unlocking and importing the drive is erased. Any drive that uses this key	No action is required. If disks need to be imported then remove the disks, wait 10 seconds, and reinsert the disks.

Table 21. Run-time issues and corrective actions (continued)

Description	Error Message	Probable Cause	Recommended Corrective Action
		which is not imported or rekeyed will be locked.	If using LKM then provide the key to unlock the disks. If using SEKM then wait for the key exchange to occur. Import the configuration after the disks are unlocked.
Error detected in the controller Level 2 cache	An error is detected in the Level-2 (L2) cache.	The controller hardware is not functioning or is going to stop functioning.	Contact your Technical Support team.
Controller flash access error	A transient error has occurred while accessing the flash device.	The controller flash device is not functioning or is going to stop functioning.	Contact your Technical Support team.
Controller flash stopped functioning	The flash device has stopped functioning.	The controller flash device is not functioning and must be replaced.	Contact your Technical Support team.
Multi-bit ECC error detected on the onboard controller memory	A multi-bit ECC error is detected in the OCM. This is a high-severity issue.	A Multi-bit ECC error is detected on the controller. Controller must be replaced.	Contact your Tech Support team to replace the controller.
Single-bit ECC error detected on the onboard controller memory	A single-bit ECC error is detected in the OCM. This is not a high-severity issue.	A Single-bit ECC error is detected.	No response action is required.
Multi-bit ECC error detection crossed warning threshold	A number of single-bit ECC errors in the OCM has crossed the Warning threshold.	A Single-bit ECC error is detected multiple times.	Continue to monitor the ECC errors. If ECC errors persists, contact your Tech Support team to replace the controller.
ECC error logging stopped	The detection of single-bit ECC errors in the OCM has been disabled.	Additional ECC errors will not be logged because too many errors are occurring and the controller does not have room to record anymore.	Contact your Technical Support team to get the controller replaced.
Controller fan stopped functioning	The fan stopped functioning because <reason>.	The fan on the adapter controller may not be functioning properly.	Check the controller and system health and replace the controller, if necessary.
Drive Initialization Error	An initialization issue is detected in the physical drive <PD Name>.	An initialization issue is detected in the physical drive during the initialization or discovery of the drive. The drive may not be in good state.	Remove or reinstall the drive. Check the health or SMART status of the physical drive. For NVMe drives, run the Recovery operation of the supported format.
Command timed out on physical drive	During backplane firmware update operation, the Command Timeout event is displayed in PERC.	The backplane may have been reset after updating the firmware, and therefore the commands sent by PERC to backplane could be timed out.	No response action is required.
An I/O delay is observed when a virtual drive state changes	An I/O delay is observed when a virtual drive state changes.	The controller has a delay time of 4 seconds for SAS/SATA and 8 seconds for NVMe drive removal processing to prevent unnecessary rebuild operations. This delay can introduce short	No response action is required.

Table 21. Run-time issues and corrective actions (continued)

Description	Error Message	Probable Cause	Recommended Corrective Action
		pause in IO when the virtual drive state changes.	
Unable to complete the Drive MAKERS Authority disable operation during the securing operation	Failed to disable the MAKERS authority on the drive.	An error occurred while securing the drive or the drive may not support MAKERS authority.	Retry the operation. If the issue persists contact support. If the issue persists, contact your Technical Support team.
Unable to complete the Drive MAKERS Authority enable operation during the erase operation	Failed to enable the MAKERS authority on the drive.	An error occurred while erasing the drive or the drive may not support MAKERS authority.	Retry the operation. If the issue persists contact support. If the issue persists, contact your Technical Support team.
Drive firmware download port lock on reset operation failed during securing operation	Failed to disable the firmware download port on the drive.	An error occurring during securing or drive may not support firmware download port control setting.	Retry the operation. If the issue persists contact support. If the issue persists, contact your Technical Support team.
Drive firmware download port unlock operation failed during erase operation	Failed to enable the firmware download port on the drive.	An error occurring during erase or Drive may not support the firmware download port control setting	Retry the operation. If the issue persists contact support. If the issue persists, contact your Technical Support team.
Firmware Update Failure during back-to-back Firmware update operations from OS.	Unknown Error: Operation failed	Controller could still be performing online activation in the background.	Ensure that you provide an interval of five minutes before attempting an back-to-back firmware update for online activation. The following message is displayed during online activation: <code>Online activation initiated to activate the downloaded firmware package.</code>

Application issues

The following table lists error messages, probable causes, and recommended response action to resolve the issues in general applications.

Table 22. General application issues and corrective actions

Description	Error Message	Probable Cause	Recommended Corrective Action
Reboot Required	Restart the server to complete the requested operation.	The operation was successful but requires a system restart to make the changes effective.	Restart the server to make the changes effective.
Reset Required	The operation is successful but reset the controller settings to make the changes effective.	The operation was successful but requires a controller reset to finish.	Reset the controller to make the changes effective.
Locked Disks Present	Unable to complete the operation because locked foreign configuration(s) is present.	The operation cannot be completed because some locked disks are present.	Remove the locked disks or cryptographically erase the locked disks to clear the security

Table 22. General application issues and corrective actions (continued)

Description	Error Message	Probable Cause	Recommended Corrective Action
			status, and then retry the operation.
Foreign Configuration	No foreign configuration is detected.	A foreign scan was run but no foreign configuration is present.	Ensure all foreign disks are present and drives are detected, then run the scan again
Low Memory	Insufficient controller memory to process the operation. Try again later.	Controller or system memory is low.	Close open applications and retry the operation. Contact Dell Support if the issue persists.
Application Compatibility	The command is unknown and not supported at this time. The command: <commandName>.	The application may not be compatible with the controller firmware.	Update application and controller firmware to the latest version.
Controller Busy	The requested command is not supported because the controller firmware initialization is not complete.	The controller is still initializing from boot or reset.	Wait some time for the controller to boot then retry the operation.
Foreign Configuration	The foreign configuration cannot be imported because the configuration is not complete.	Some disks are missing from the foreign configuration.	Ensure all foreign disks are present and drives are detected, then run the scan again.
Drive Security	The requested operation cannot be completed because the controller already has the security key.	The operation is not supported when the controller is in either the LKM or EKM mode.	Disable the drive security mode on the controller and retry the operation.
Drive Security	The requested operation cannot be completed because the security key is not present.	Operation is not supported when the controller is not in the LKM or EKM mode.	Disable the drive security mode on the controller and retry the operation.
Drive Security	The requested operation cannot be completed because the security key is invalid.	An incorrect Security Key or Key ID was entered when changing the key.	Enter the correct key and retry the operation.
Application Compatibility	The requested operation cannot be performed because of an internal error.	The application may not be compatible with the controller firmware.	Update application and controller firmware to the latest version.
Snapdump Error	The requested operation cannot be performed because of the existing snapdump. Clear all the snapdump and reinitiate the operation.	Some operations cannot be performed when there are debug log snapdumps present on the controller.	Download all snapdumps from the controller then clear the snapdumps and retry the operation.
Snapdump Error	The requested operation cannot be performed because no snapdump is present.	No snapdumps are present when trying to download debug log.	Generate a new on-demand snapdump.
Snapdump Error	The requested operation cannot be performed because the on-demand snapdump is not allowed.	The controller does not support snapdump generation or is in a state where a snapdump cannot be generated such as bad DDR or is in a fault condition.	Ensure the controller supports snapdump functionality. Check the controller's DDR is healthy and that the controller is not in a fault state.
Snapdump Error	The requested operation cannot be performed because there is an ongoing on-demand snapdump.	There is an ongoing snapdump being collected. Note that On-demand snapdump can only be collected once every 10 minutes.	Retry the operation after some time.
Hardware Error	The requested operation cannot be completed because of a hardware error. The extended Status code: <statusCode>.	The controller may be in a faulty state or the hardware may not be fully functional.	Retry the operation after some time. If the issue persists, contact the Dell Technical Support team.

Table 22. General application issues and corrective actions (continued)

Description	Error Message	Probable Cause	Recommended Corrective Action
Firmware Error	The requested operation cannot be completed because of a controller firmware error. The extended Status code: <statusCode>.	The controller may be in a faulty state or the has an internal error.	Retry the operation after some time. If the issue persists, contact the Dell Technical Support team.
Invalid ID	The requested operation cannot be completed because the ID is invalid. The extended Status code: <statusCode>.	The application is no longer in sync with the firmware or the device is unavailable.	Retry the operation after some time. If the issue persists, contact the Dell Technical Support team.
Device not Found	The requested operation cannot be completed because the device is not detected. The extended Status code: <statusCode>.	The requested device is not available.	Ensure that the correct device ID is selected when running the commands. If the issue persists, contact the Dell Technical Support team.
Invalid Sequence Number	Check the configuration and retry the operation. The extended Status code: <statusCode>.	The application is no longer in sync with the firmware.	Retry the operation after some time. If the issue persists, contact the Dell Technical Support team.
Invalid Argument	The requested command has invalid arguments. The extended Status code: <statusCode>.	The application is no longer compatible with the controller firmware or no longer in sync with the firmware.	Update the application and controller firmware to the latest version and retry the operation after some time. If the issue persists, contact the Dell Technical Support team.
Command Not supported for drive	The requested command is not supported for the drive. The extended Status code: <statusCode>.	The VD may be in a state where the specific operation is not supported or the command is not supported at all.	Ensure that the controller and drive support the feature. For performing erase operations on PDs, the drives must be in Unconfigured state.
Command Not supported for VD	The requested command is not supported for the virtual drive. The extended Status code: <statusCode>.	The VD may be in a state where the specific operation is not supported or the command is not supported at all.	Ensure that the controller and VD support the feature.
Command Not supported for Controller	The requested command is not supported. The extended Status code: <statusCode>.	The Controller or Configuration may be in a state where the specific operation is not supported or the command is not supported at all.	Confirm Controller supports the feature. Ensure that the controller and VDs are healthy state.
Controller Busy	The requested command cannot be completed because the maximum limit is exceeded. The extended Status code: <statusCode>.	Too many commands are run on the controller.	Wait some time for some commands to complete then retry the operation.
Drive Mixing Violation	Mixing of the drive type is not supported. Add the same type of drives to perform the operation. The extended Status code: <statusCode>.	Mixing of some physical disk properties are not supported (Sector Size, Media Type, or Protocol).	Ensure the same type of drives are being used for the operation.
Force Option Required	This command can be processed only with the Force option.	The command may be erasing data or causing a short performance degradation.	The Force parameter is required for any operation that may erase data or cause performance degradation. Select or add the option and retry running the command.
Configuration Present	This command can be processed only when no configuration is present.	The operation cannot be run while a configuration is present.	Clear the configuration from the controller or remove all

Table 22. General application issues and corrective actions (continued)

Description	Error Message	Probable Cause	Recommended Corrective Action
			the configured devices from the system.
Controller not responding	The controller firmware is not responding.	The controller may be busy processing other commands or has an issue.	Wait some time and try the operation again. If the issue persists then try restarting the server.
PCI Error	The requested operation cannot be performed because of a PCI error.	A PCIe error has occurred.	N/A
Flash error	The requested operation cannot be performed because of a flash error.	An error occurred on the hardware flash.	Wait some time then retry the operation. If the issue persists, contact the Dell Technical Support team.
Drive Error	The requested operation cannot be performed because of an error is detected on the drive.	An error occurred on the drive.	Wait some time then retry the operation. If the issue persists, contact the Dell Technical Support team.
Storage space on the drive is too less	The requested operation cannot be performed because storage space of the drive is too less.	The replaced drive has too less storage capacity for the array.	Replace the drive with a drive that has high storage capacity.
Drive does not support security	The requested operation cannot be performed because the drive is not security-capable.	The selected or replaced drive is not security-capable.	Replace the drive with a security-capable drive.
Drive security protocol mismatch	The requested operation cannot be performed because the security type of the drive is inappropriate.	The selected or replaced drive does not support the compatible security protocol.	Replace the drive with a drive that is compatible with the security protocol.
Drive Type Mismatch	The requested operation cannot be performed because the drive type is incorrect.	The selected or replaced drive does not match with the drives already in the array.	Ensure that the drive type matches either SAS, SATA, or NVMe; 512n or 4kn; or HDD or SSD.
Drive Sector Type mismatch	The requested operation cannot be performed because the block size of the drive does not match with other drives in the array.	The block size of the selected or replaced drive does not match with the drives already in the array	Ensure that the sector size of the drive matches 512n or 4kn.
Security not enabled	The requested operation cannot be performed because the security feature of the drive is not enabled.	The security mode of the controller is not enabled.	Enable security on the controller and retry the operation.
Wrong Enclosure	The command cannot be run because the drive is not a part of this enclosure.	An incorrect enclosure ID is entered.	Enter a correct enclosure ID in the command and retry the operation.
No SEKM capable Agent Detected	The requested operation cannot be performed because the Security External Key Management (SEKM) server is not detected.	iDRAC may not be compatible with SEKM or does not support the SEKM feature.	Update iDRAC and ensure that SEKM features are supported.
Command Not supported for Raid Level	The requested operation cannot be performed for the current RAID level.	The operation may not be supported for RAID levels such as running BGI or CC on a RAID 0 array.	Ensure that the operation is supported for the RAID array.
VD Wrong state	The requested operation cannot be performed because the VD is in a wrong state.	The VD may be in a degraded or offline state where the operation is unsupported.	Ensure the VD is in a healthy state.

Table 22. General application issues and corrective actions (continued)

Description	Error Message	Probable Cause	Recommended Corrective Action
Too many retries	Unable to run the command because the maximum number of retries is exceeded.	The command is run for more than the maximum stipulated number of attempts.	Wait for some time and retry the operation. If the issue persists, contact the Dell Technical Support team.
Preserved Cache Present	The requested operation cannot be performed because the controller has data in cache for offline or missing virtual drive(s).	The operation cannot be performed because there is preserved cache present on the controller.	Ensure all drives are detected, and then import any foreign drives. Also, clear the preserved cache.
Feature Not Supported	The requested operation cannot be performed because the feature is not supported.	The feature is not supported on the controller.	No response action is required because the feature is not supported.
Secure Devices Present	The requested operation cannot be completed because there are one or more secure drives present on this controller.	The operation cannot be performed because there are secured drives present on the controller.	By using the cryptographic erase operation erase the secured drives or remove the drives from the server, and then retry the operation.
Operation disabled	The requested operation cannot be performed because the operation is currently disabled.	The operation cannot be performed because the feature is disabled or turned off.	Re-enable the feature and retry the operation.
Operation in Progress	The requested operation cannot be performed because some operation is currently in progress.	The operation cannot be performed because there is already another operation in progress.	Wait for the current operation to complete or stop the current operation.
Controller in SafeMode	The requested command is not supported because the controller is running in safe mode.	The controller is running in safe mode because of some internal issues.	Check the health of the controller and configuration and correct any errors
Controller Faulted	An issue is detected with the controller firmware.	An issue is detected in the controller which can result in a auto-reset or snapdump operation.	Save the snapdump generated by the controller and Contact the Dell Technical Support team.
Drive Security State error	The requested operation cannot be performed because of the current SED state of the drive.	The operation is not supported because the drive is in an unsupported security state. Drive may be secured when it should be unsecured.	Check the security state of the drive. Perform cryptographic erase or secure operation as needed.
Auto-Configuration Enabled	The requested operation cannot be completed because of a secured auto-configure setting.	The auto-configuration option is enabled.	Change the auto-configuration setting to the default (Unconfigured) setting then retry the operation.

Memory errors

Memory errors can corrupt cached data, so the controllers are designed to detect and attempt to recover from the memory errors. Single-bit memory errors can be handled by the controller and do not disrupt normal operation. A notification is sent if the number of single-bit errors exceeds a threshold value.

Multi-bit errors are more serious because the effected memory data cannot be automatically recovered. The following are the actions that occur in the case of multi-bit errors:

- When a multi-bit error is detected, the controller stops responding and becomes inoperative. You must restart the controller.
- If a multi-bit error occurs while accessing data in the cache when the controller is started with dirty cache, the controller discards the cache contents. The controller generates a Warning message to the system console to indicate that the cache was discarded, and then generates an event.

- In all cases, the controller logs an event to the controller's internal event log and a message is displayed during POST indicating that a multi-bit error has occurred.

 **NOTE:** In case of a multi-bit error, contact [Contact Technical Support](#).

Preserved Cache state

The controller preserves the dirty cache from a virtual disk if the virtual disk goes offline or is deleted because of missing physical disks. This preserved dirty cache is called preserved cache and is preserved until you import the virtual disk or discard the cache.

1. Import the virtual disk—Power off the system, re-insert the virtual disk and restore the system power. Use the **HII Configuration Utility** to import the foreign configuration.
2. Discard the preserved cache—See [Discard preserved cache](#).

 **NOTE:** It is recommended to clear the preserved cache before reboot using any of the virtual disks present on the controller.

Security key errors

Secured foreign import errors

A foreign configuration is a RAID configuration that already exists on a replacement physical disk that you install in a system. A secured foreign configuration is a RAID configuration that was created under a different security key.

There are two scenarios in which a secured foreign import fails:

- **The security key authentication fails**—A virtual disk secured with a security key different from the current controller security key cannot be imported without authentication of the original security key used to secure them. Supply the correct security key to import the secured foreign configuration. If you have lost or forgotten the security key, the secured foreign disks remain locked (inaccessible) until the appropriate security key is entered or disks are erased.
- **The secured virtual disk is in an offline state after supplying the correct security key**—You must check to determine why the virtual disk failed and correct the problem.

Failure to select or configure non Self-Encrypting Disks non-SED

A virtual disk can be either secured or unsecured depending on how it was configured when created. In order to create a secured virtual disk, the controller must have a security key present and must contain SEDs only. In order to select/configure non-SED, you must create an unsecured virtual disk. You can create an unsecured virtual disk even if there is a security key. Select the **Secure VD** option as **No** in the **Create New VD** menu. For steps on how to create an unsecured virtual disk, see [Create virtual disk and configure virtual disk parameters](#).

Failure to delete security key

A security key is used to lock or unlock access to a security-enabled component. This key is not utilized in the actual encryption of data. If a security key is present, both secured and unsecured virtual disks may exist.

To delete the security key, you must have a previously established security key present on the controller and there cannot be any configured secured disks. If there are configured secured virtual disks, remove or delete them.

Failure of Cryptographic Erase on encryption-capable physical disks

Cryptographic Erase is the process of securely erasing all data permanently on an encryption-capable physical disk and resetting the security attributes. It is used in scenarios such as deleting a foreign configuration in the event of a forgotten or lost security key, or unlocking a disk that had been previously locked.

 **NOTE:** PERC12 does not support unlocking disks from past generations or from third-party secured configurations.

You can perform Cryptographic Erase only on encryption-capable disks that are not hot spares and not configured as non-RAID or virtual disks. Ensure that the conditions are met. See [Cryptographic erase](#).

General issues

PERC goes into Safe and Unresponsive state

- Issue** After performing the rekey operation on the volume of a drive where the iLKM feature is enabled, and then restarting the server, the Red Screen of Death (RSOD) issue is observed. The PERC state becomes Safe Mode or Unresponsive.
- Corrective Action:** Remove the locked drives from the server and restart the server. After restarting, install the locked drives again into the server.

Enclosure state is indicated as faulty instead of need attention

- Issue** When a drive that cannot be accessed is removed, the enclosure state is indicated as faulty instead of Need Attention.
- Corrective Action:** Expected behavior. No action is required.

perccli2 commands and OS DUPs must have the controller host driver 8.0.0.69.0 for proper functioning

- Issue** The perccli2 commands and operating system DUPs require the controller host driver minimum version of 8.0.0.69.0 for proper functioning. They do not work on inbox drivers earlier than 8.0.0.69.0.
- Corrective Action:** Expected behavior. No action is required.

PERC card has yellow bang in Windows operating system device manager

- Issue:** The device is displayed in **Device Manager** but has a yellow exclamation mark.
- Corrective Action:** Reinstall the driver. For more information on reinstalling drivers, see [Driver support for PERC H965i Adapter, PERC H965i Front, PERC H965i MX, and PERC H965e Adapter](#).

PERC card not seen in operating systems

- Issue:** The device does not appear in the **Device Manager**.
- Corrective Action:** Power off the system and reseal the controller. For more information, see [Install and remove a PERC 12 card](#).

Unlocked drives if not imported appear as locked

- Issue:** If drives are unlocked but not imported, the drives appear as locked even after the controller is reset.
- Corrective Action:** Remove the drives from the system.

Physical disk issues

Lifecycle log indicates drive number as unknown

- Issue** After performing a cryptographic operation, Lifecycle log data displays the serial number of a drive as unknown. Also, PPID, product ID, revision, and serial number data are not displayed on the Lifecycle Controller user interface.
- Corrective Action:** Expected behavior. Perform a Cold Reboot operation on the server.

An NVMe drive is not indicated as unsupported when EEDP is enabled

- Issue** An NVMe drive on which the EEDP feature is enabled is not indicated as unsupported when the PI type is set to 1, 2, or 3.
- Corrective Action:** Expected behavior. No action is required. Dell does not support NVMe drives on which the EEDP feature is enabled.

Unexpected generation of incorrect physical AEN

- Issue** After initializing an NVMe drive, an unexpected and incorrect physical AEN is generated. Enclosure does not function correctly.
- Corrective Action:** Remove the NVMe drive that has no namespace in it.

Unconfigured drives are not listed in the OS or Hardware Diagnostics

- Issue** Unconfigured drives (ready-state drives) are not listed in the OS or Hardware Diagnostics of Lifecycle Controller.
- Corrective Action:** Expected behavior. No action is required.

Physical disk in failed state

- Issue:** One of the physical disks in the disk array is in the failed state.
- Corrective Action:** Update the PERC cards to the latest firmware available on the support site and replace the drive.

The PDR6 Error and Event message is saved in the log data when a new drive is inserted

- Issue:** The PDR6 Error and Event message is saved in the log data when a new drive is inserted.
- Probable Cause:** When the auto rebuild feature is enabled and the drive is inserted, then the drive status becomes Offline and then Rebuilding.
- Corrective Action:** Expected behavior. No action is required because this behavior does not imply that the drive has permanently become offline.

Unable to rebuild a fault tolerant virtual disk

Issue:	Cannot rebuild a fault tolerant virtual disk. For more information, see the alert log for virtual disks.
Probable Cause:	The replacement disk is too small or not compatible with the virtual disk.
Corrective Action:	Replace the failed disk with a compatible good physical disk with equal or greater capacity.

An NVMe drive is indicated as faulty when moving the drive to new backplane using hot-insertion

Issue:	When moving an NVMe drive to a different backplane slot, if the NVMe drive is hot-removed and hot-inserted in a different slot too quickly, then the drive may be indicated as faulty.
Probable Cause:	NVMe drive is moved between slots too quickly.
Corrective Action:	Wait at least 8 seconds after a hot-pull event before re-inserting the drive in an alternate slot.

Fatal error or data corruption reported

Issue:	Fatal error(s) or data corruption(s) are reported when accessing virtual disks.
Corrective Action:	Contact your Technical Support team.

Multiple disks are inaccessible

Issue:	Multiple disks are simultaneously inaccessible.
Probable Cause:	Multiple physical disk errors in a single array typically indicate a failure in cabling or connection and could involve the loss of data.
Corrective Action:	You can recover the virtual disk after multiple physical disks become simultaneously inaccessible. Perform the following steps to recover the virtual disk:

 **CAUTION:** Follow the safety precautions to prevent electrostatic discharge.

1. Turn off the system, check cable connections, and reseal physical disks.
2. Ensure that all the disks are present in the enclosure.
3. Turn on the system and enter the **HII Configuration Utility**.
4. Import the foreign configuration.
5. Press <F> at the prompt to import the configuration, or press <C> to enter the **HII Configuration Utility** and either import or clear the foreign configuration.

If the virtual disk is redundant and transitioned to **Degraded** state before going **Offline**, a rebuild operation starts automatically after the configuration is imported. If the virtual disk has gone directly to the **Offline** state due to a cable pull or power loss situation, the virtual disk is imported in its **Optimal** state without a rebuild occurring.

 **NOTE:** You can use the **HII Configuration Utility** or Dell OpenManage storage management application to perform a manual rebuild of multiple physical disks.

Rebuilding data for a failed physical disk

Issue:	Rebuilding data for a physical disk that is in a failed state.
Probable Cause:	Physical disk is failed or removed.
Corrective Action:	If you have configured hot-spares, the PERC card automatically tries to use one of the hot-spares to rebuild a physical disk that is in a failed state. Manual rebuild is necessary if no hot-spares with enough

capacity to rebuild the failed physical disks are available. You must insert a physical disk with enough storage in the subsystem before rebuilding the physical disk.

NOTE: You can use the **HII Configuration Utility** or Dell OpenManage storage management application to perform a manual rebuild of an individual physical disk.

Virtual disk fails during rebuild using a global hot-spare

Issue: A virtual disk fails during rebuild while using a global hot spare.

Probable Cause: One or more disks in the virtual disks fails or is disconnected while the rebuild is in progress.

Corrective Action: No action is required. The global hot spare reverts to **Hot spare** state and the virtual disk is in **Failed** state.

Dedicated hot-spare disk fails during rebuild

Issue: A hot-spare disk fails during rebuild while using a dedicated hot spare.

Probable Cause: The dedicated hot-spare assigned to the virtual disk fails or is disconnected while the rebuild is in progress.

Corrective Action: If there is a global hot spare available with enough capacity, rebuild will automatically start on the global hot spare. Where there is no hot spare present, you must insert a physical disk with enough capacity into the system before performing a rebuild.

Virtual disk fails rebuild using a dedicated hot spare

Issue: A virtual disk fails during rebuild while using a dedicated hot spare.

Probable Cause: One or more disks in the virtual disks fails or is disconnected while the rebuild is in progress.

Corrective Action: No action is required. The dedicated hot spare is in **hot spare** state and converted to global hot spare if there is any other virtual disk that is supported, otherwise the dedicated hot spare reverts to **Ready** state and the virtual drive is in **Failed** state.

Physical disk takes a long time to rebuild

Issue: A physical disk is taking longer than expected to rebuild.

Description: A physical disk takes longer to rebuild when under high I/O stress. There is only one rebuild I/O operation for every five host I/O operations.

Corrective Action: If possible, reduce I/O stress on the physical disk or increase the value of rebuild rate controller parameter.

Drive removal and insertion in the same slot generates a foreign configuration event

Issue: When a drive which is part of a virtual disk is removed and reinserted into the same slot the drive goes through a transient state of being foreign for a short period of time before rebuilding.

Description: This transient state could be reported as an event in management applications as **A foreign configuration was detected on RAID Controller is SL x**, where x is the slot of the RAID controller.

Corrective Action: No action is required on the foreign configuration state of the drive as it is transient and the controller handles the event automatically.

SMART errors

SMART monitors the internal performance of all motors, heads, and physical disk electronics and detects predictable physical disk failures.

i **NOTE:** For information about SMART errors' reports that could indicate hardware failure, see the *Dell OpenManage Storage Management User's Guide* available at [OpenManage Manuals](#).

SMART error detected on a non-RAID disk

- Issue:** A SMART error is detected on a non-RAID disk.
- Corrective Action:** Perform the following steps:
1. Back up your data.
 2. Replace the affected physical disk with a new physical disk of equal or higher capacity.
 3. Restore from the backup.

SMART error detected on a physical disk in a non-redundant virtual disk

- Issue:** A SMART error is detected on a physical disk in a non-redundant virtual disk.
- Corrective Action:** Perform the following steps:
1. Back up your data.
 2. Use **Replace Member** to replace the disk manually.
- i** **NOTE:** For more information about the **Replace Member** feature, see [Configure hot-spare drives](#).
3. Replace the affected physical disk with a new physical disk of equal or higher capacity.
 4. Restore from the backup.

SMART error detected on a physical disk in a redundant virtual disk

- Issue:** A SMART error is detected on a physical disk in a redundant virtual disk.
- Corrective Action:** Perform the following steps:
1. Back up your data.
 2. Force the physical disk offline.
- i** **NOTE:** If a hot-spare is present, the rebuild starts with the hot-spare after the disk is forced offline.
3. Replace the disk with a new physical disk of equal or higher capacity.
 4. Perform the **Replace Member** operation.
- i** **NOTE:** The **Replace Member** operation allows you to copy data from a source physical disk of a virtual disk to a target physical disk that is not a part of the virtual disk. For more information about the **Replace Member** feature, see [Configure hot-spare drives](#).

Replace member errors

i **NOTE:** For more information about the **Replace Member** features, see [Configure hot-spare drives](#).

Source disk fails during replace member operation

- Issue:** The source disk fails during the **Replace Member** operation and the **Replace Member** operation stops due to the source physical disk error.
- Probable Cause:** Physical disk failure or physical disk is removed or disconnected.
- Corrective Action:** No action required. If the virtual disk can tolerate disk failure, and the source data is available from other disks in the virtual disk, the rebuild begins automatically on the target disk, using the data from the other disks, if the virtual disk cannot tolerate the failure, the virtual disk goes to offline state and the replace member operation is stopped.

Target disk fails during replace member operation

- Issue:** The target disk failure reported during the **Replace Member** operation, and the **Replace Member** operation stops.
- Probable Cause:** Physical disk failure or physical disk is removed or disconnected.
- Corrective Action:** It is recommended that you replace or check the target drive, and restart the **Replace Member** operation or perform the operation on a different target drive.

A member disk failure is reported in the virtual disk which undergoes replace member operation

- Issue:** The source and the target drive which is part of **Replace Member** operation are online, while a different drive which is a member of the virtual drive reports a failure.
- Probable Cause:** Physical disk failure or physical disk is removed or disconnected.
- Corrective Action:** A rebuild starts if there any hot-spares configured or you may replace the failed drive. The **Replace Member** operation continues as far as the source virtual disk can tolerate the drive failure. If the source virtual disk fails, the **Replace Member** is stopped, otherwise the virtual disk continues to be in degraded state.

Linux operating system errors

Virtual disk policy is assumed as write-through

- Error:** `<Date:Time> <HostName> kernel: sdb: asking for cache data failed<Date:Time> <HostName> kernel: sdb: assuming drive cache: write through`
- Corrective Action:** The error message is displayed when the Linux Small Computer System Interface (SCSI) mid-layer asks for physical disk cache settings. The controller firmware manages the virtual disk cache settings on a per controller and a per virtual disk basis, so the firmware does not respond to this command. The Linux SCSI mid-layer assumes that the virtual disk's cache policy is **Write-Through**. SDB is the device node for a virtual disk. This value changes for each virtual disk. For more information about **Write-Through** cache, see [Virtual disk write cache and read-ahead policy](#). Except for this message, there is no effect of this behavior on normal operation. The cache policy of the virtual disk and the I/O throughput are not affected by this message. The cache policy settings for the PERC SAS RAID system remain unchanged.

Unable to register SCSI device

Error: smartd[smartd[2338] Device: /dev/sda, Bad IEC (SMART) mode page, err=-5, skip device smartd[2338] Unable to register SCSI device /dev/sda at line 1 of file /etc/smartd.conf.

Corrective Action: This is a known issue. An unsupported command is entered through the user application. User applications attempt to direct Command Descriptor Blocks to RAID volumes. The error message does not affect the feature functionality. The Mode Sense/Select command is supported by firmware on the controller. However, the Linux kernel **daemon** issues the command to the virtual disk instead of to the driver **IOCTL** node. This action is not supported.

Drive indicator codes

The LEDs on the drive carrier indicates the state of each drive. Each drive carrier has two LEDs: an activity LED (green) and a status LED (bicolor, green/amber). The activity LED blinks whenever the drive is accessed.



Figure 17. Drive indicators

1. Drive activity LED indicator
2. Drive status LED indicator
3. Drive capacity label

If the drive is in the Advanced Host Controller Interface (AHCI) mode, the status LED indicator does not power on. Drive status indicator behavior is managed by Storage Spaces Direct. Not all drive status indicators may be used.

NOTE: For information about the Drive LED indicator codes in an MD24XX series enclosure, see the Dell PowerVault MD24XX Direct-Attach Storage for PowerEdge Servers Owner's Manual available on the support site.

Table 23. Drive indicator codes

Drive status indicator code	Condition
Blinks green twice per second	The drive is being identified or preparing for removal
Off	The drive is ready for removal NOTE: The drive status indicator remains off until all drives are initialized after the system is powered on. Drives are not ready for removal during this time.
Blinks green, amber, and then powers off	There is an expected drive failure
Blinks amber four times per second	The drive has failed or faulted or not detected.
Blinks green slowly	The drive is rebuilding
Solid green	The drive is online

Table 23. Drive indicator codes (continued)

Drive status indicator code	Condition
Blinks green for three seconds, amber for three seconds, and then powers off after six seconds	The rebuild has stopped

The AutoSecure feature is not supported

Issue The AutoSecure feature of iDRAC is not supported by PERC 12 and HBA465 controllers.

Corrective Action: Expected behavior. No action is required.

HII error messages

Unhealthy Status of the drivers

Error: One or more boot driver(s) have reported issues. Check the Driver Health Menu in Boot Manager for details.

Probable Cause: This message might indicate that the cables are not connected, the disks might be missing, or the UEFI driver might require configuration changes.

Corrective Action:

1. Check if the cables are connected properly, or replace missing hard drives, if any and then restart the system.
2. Press any key to load the driver health manager to display the configurations. The Driver Health Manager displays the driver(s), which requires configuration.
3. Alternately, if the UEFI driver requires configuration, press any key to load the Configuration Utility.

Rebuilding a drive during full initialization

Issue: Automatic rebuild of drives is disabled for virtual disk during full initialization.

Corrective Action: After full initialization the drive will automatically start its rebuild on its corresponding virtual disk.

Appendix—RAID description

RAID is a group of independent physical disks that provides high performance by increasing the number of disks used for saving and accessing data.

⚠ CAUTION: In the event of a physical disk failure, a RAID 0 virtual disk fails, resulting in data loss.

A RAID disk subsystem offers the following benefits:

- Improved I/O performance and data availability.
- Improved data throughput because several disks are accessed simultaneously. The physical disk group appears either as a single storage unit or multiple logical units to the host system.
- Improved data storage availability and fault tolerance. Data loss caused by a physical disk failure can be recovered by rebuilding missing data from the remaining physical disks containing data or parity.

Topics:

- [Summary of RAID levels](#)
- [RAID 10 configuration](#)
- [RAID terminology](#)

Summary of RAID levels

Following is a list of the RAID levels supported by the PERC 12 series of cards:

- RAID 0 uses disk striping to provide high data throughput, especially for large files in an environment that requires no data redundancy.
- RAID 1 uses disk mirroring so that data written to one physical disk is simultaneously written to another physical disk. RAID 1 is good for small databases or other applications that require small capacity and complete data redundancy.
- RAID 5 uses disk striping and parity data across all physical disks (distributed parity) to provide high data throughput and data redundancy, especially for small random access.
- RAID 6 is an extension of RAID 5 and uses an additional parity block. RAID 6 uses block-level striping with two parity blocks distributed across all member disks. RAID 6 provides protection against double disk failures, and failures while a single disk is rebuilding. If you are using only one array, deploying RAID 6 is more effective than deploying a hot spare disk.
- RAID 10 is a combination of RAID 0 and RAID 1, uses disk striping across mirrored disks. It provides high data throughput and complete data redundancy.
- RAID 50 is a combination of RAID 0 and RAID 5 where a RAID 0 array is striped across RAID 5 elements. RAID 50 requires at least six disks.
- RAID 60 is a combination of RAID 0 and RAID 6 where a RAID 0 array is striped across RAID 6 elements. RAID 60 requires at least eight disks.

The following table lists the minimum and maximum disks supported on each RAID levels.

Table 24. Minimum and maximum disks supported on each RAID levels

RAID Level	Minimum disk	Maximum disk
0	1	32
1	2	2
5	3	32
6	4	32
10	4	240
50	6	240
60	8	240

RAID 10 configuration

Any RAID 10 volume that has more than 32 drives require spanning. Each span can contain up to 32 drives. Drives must be distributed evenly across all the spans with each span containing an even number of drives.

NOTE: Spans in a RAID 10 volume are only supported if spans are even. Uneven spanned RAID 10 cannot be imported from previous controller generations.

The following table shows the RAID 10 configurations.

Table 25. RAID 10 configurations

Disk or span count	RAID 10 capable	Disk or span count	RAID 10 capable	Disk or span count	RAID 10 capable	Disk or span count	RAID 10 capable
4 (1)	Yes	64 (2)	Yes	124	No	184	No
6 (1)	Yes	66 (3)	Yes	126 (7)	Yes	186	No
8 (1)	Yes	68	No	128 (4)	Yes	188	No
10 (1)	Yes	70 (5)	Yes	130 (5)	Yes	190	No
12 (1)	Yes	72 (3)	Yes	132 (6)	Yes	192 (6)	Yes
14 (1)	Yes	74	No	134	No	194	No
16 (1)	Yes	76	No	136	No	196 (7)	Yes
18 (1)	Yes	78 (3)	Yes	138	No	198	No
20 (1)	Yes	80 (4)	Yes	140 (5)	Yes	200	No
22 (1)	Yes	82	No	142	No	202	No
24 (1)	Yes	84 (6)	Yes	144	Yes	204	No
26 (1)	Yes	86	No	146	No	206	No
28 (1)	Yes	88 (4)	Yes	148	No	208 (8)	Yes
30 (1)	Yes	90 (3)	Yes	150 (5)	Yes	210 (7)	Yes
32 (1)	Yes	92	No	152	No	212	No
34	No	94	No	154 (7)	Yes	214	No
36 (2)	Yes	96 (3)	Yes	156 (6)	Yes	216	No
38	No	98 (7)	Yes	158	No	218	No
40 (2)	Yes	100 (5)	Yes	160 (5)	Yes	220	No
42 (2)	Yes	102	No	162	No	222	No
44 (2)	Yes	104 (4)	Yes	164	No	224 (8)	Yes
46	No	106	No	166	No	226	No
48 (2)	Yes	108 (6)	Yes	168 (6)	Yes	228	No
50 (2)	Yes	110 (5)	Yes	170	No	230	No
52 (2)	Yes	112 (4)	Yes	172	No	232	No
54 (2)	Yes	114	No	174	No	234	No
56 (2)	Yes	116	No	176 (8)	Yes	236	No
58	No	118	No	178	No	238	No
60 (2)	Yes	120 (4)	Yes	180 (6)	Yes	240 (8)	Yes
62	No	122	No	182 (7)	Yes	-	-

RAID terminology

Disk striping

Disk striping allows you to write data across multiple physical disks instead of just one physical disk. Disk striping involves partitioning each physical disk storage space in stripes of the following sizes: 64 KB, 128 KB, 256 KB, 512 KB, and 1 MB. The stripes are interleaved in a repeated sequential manner. The part of the stripe on a single physical disk is called a stripe element.

For example, in a four-disk system using only disk striping (used in RAID 0), segment 1 is written to disk 1, segment 2 is written to disk 2, and so on. Disk striping enhances performance because multiple physical disks are accessed simultaneously, but disk striping does not provide data redundancy.




Figure 18. Example of disk striping (RAID 0)

Disk mirroring

With mirroring (used in RAID 1), data written to one disk is simultaneously written to another disk. If one disk fails, the contents of the other disk can be used to run the system and rebuild the failed physical disk. The primary advantage of disk mirroring is that it provides complete data redundancy. Both disks contain the same data at all times. Either of the physical disks can act as the operational physical disk.

Disk mirroring provides complete redundancy, but is an expensive option because each physical disk in the system must be duplicated.

NOTE: Mirrored physical disks improve read performance by read load balance.




Figure 19. Example of Disk Mirroring (RAID 1)

Spanned RAID levels

Spanning is a term used to describe the way in which RAID levels 10, 50, and 60 are constructed from multiple sets of basic, or simple RAID levels. For example, a RAID 10 has multiple sets of RAID 1 arrays where each RAID 1 set is considered a span. Data is then striped (RAID 0) across the RAID 1 spans to create a RAID 10 virtual disk. Similarly, RAID 50 and RAID 60 combine multiple sets of RAID 5 or RAID 6 respectively with striping.

Parity data

Parity data is redundant data that is generated to provide fault tolerance within certain RAID levels. In the event of a disk failure, the parity data can be used by the controller to regenerate user data. Parity data is present for RAID 5, 6, 50, and 60.

The parity data is distributed across all the physical disks in the system. If a single physical disk fails, it can be rebuilt from the parity and the data on the remaining physical disks. RAID level 5 combines distributed parity with disk striping. Parity provides redundancy for one physical disk failure without duplicating the contents of the entire physical disks.

RAID 6 combines dual distributed parity with disk striping. This level of parity allows for two disk failures without duplicating the contents of entire physical disks.




Figure 20. Example of Distributed Parity (RAID 5)

NOTE: Parity is distributed across multiple physical disks in the disk group.




Figure 21. Example of Dual Distributed Parity (RAID 6)

NOTE: Parity is distributed across all disks in the array.

Getting help

Topics:

- [Recycling or End-of-Life service information](#)
- [Contacting Dell](#)
- [Locating the Express Service Code and Service Tag](#)

Recycling or End-of-Life service information

Take back and recycling services are offered for this product in certain countries. If you want to dispose of system components, visit the Dell Recycle Worldwide website and select the relevant country.

Contacting Dell

Dell provides online and telephone based support and service options. If you do not have an active internet connection, you can find Dell contact information on your purchase invoice, packing slip, bill or Dell product catalog. The availability of services varies depending on the country and product, and some services may not be available in your area. To contact Dell for sales, technical assistance, or customer service issues:

Steps

1. Visit the support site.
2. Select your country from the drop-down menu on the lower right corner of the page.
3. For customized support:
 - a. Enter the system Service Tag in the **Enter a Service Tag, Serial Number, Service Request, Model, or Keyword** field.
 - b. Click **Submit**.
The support page that lists the various support categories is displayed.
4. For general support:
 - a. Select your product category.
 - b. Select your product segment.
 - c. Select your product.
The support page that lists the various support categories is displayed.
5. For contact details of Dell Global Technical Support:
 - a. Click **Contact Technical Support**.
 - b. The **Contact Technical Support** page is displayed with details to call, chat, or e-mail the Dell Global Technical Support team.

Locating the Express Service Code and Service Tag

The unique Express Service Code and Service Tag is used to identify the system.

The information tag is located on the front of the system rear of the system that includes system information such as Service Tag, Express Service Code, Manufacture date, NIC, MAC address, QRL label, and so on. If you have opted for the secure default access to iDRAC, the Information tag also contains the iDRAC secure default password. If you have opted for iDRAC Quick Sync 2, the Information tag also contains the OpenManage Mobile (OMM) label, where administrators can configure, monitor, and troubleshoot the PowerEdge servers.




Figure 22. Locating the Express Service Code and Service tag

1. Information tag (front view)
2. Information tag (back view)
3. OpenManage Mobile (OMM) label
4. iDRAC MAC address and iDRAC secure password label
5. Service Tag, Express Service Code, QRL label

The Mini Enterprise Service Tag (MEST) label is located on the rear of the system that includes Service Tag (ST), Express Service Code (Exp Svc Code), and Manufacture Date (Mfg. Date). The Exp Svc Code is used by Dell to route support calls to the appropriate personnel.

Alternatively, the Service Tag information is located on a label on left wall of the chassis.

Documentation resources

This section provides information about the documentation resources for your system.

To view the document that is listed in the documentation resources table:

- From the Dell support site:
 1. Click the documentation link that is provided in the Location column in the table.
 2. Click the required product or product version.

 **NOTE:** To locate the product name and model, see the front of your system.

3. On the Product Support page, click **Manuals & documents**.
- Using search engines:
 - Type the name and version of the document in the search box.

Table 26. Additional documentation resources for your system

Task	Document	Location
Setting up your system	For more information about installing and securing the system into a rack, see the Rail Installation Guide included with your rail solution. For information about setting up your system, see the <i>Getting Started Guide</i> document that is shipped with your system.	PowerEdge Server Manuals
Configuring your system	For information about the iDRAC features, configuring and logging in to iDRAC, and managing your system remotely, see the Integrated Dell Remote Access Controller User's Guide. For information about understanding Remote Access Controller Admin (RACADM) subcommands and supported RACADM interfaces, see the RACADM CLI Guide for iDRAC. For information about Redfish and its protocol, supported schema, and Redfish Eventing implemented in iDRAC, see the Redfish API Guide. For information about iDRAC property database group and object descriptions, see the Attribute Registry Guide. For information about Intel QuickAssist Technology, see the Integrated Dell Remote Access Controller User's Guide.	PowerEdge Server Manuals
	For information about earlier versions of the iDRAC documents. To identify the version of iDRAC available on your system, on the iDRAC web interface, click	iDRAC Manuals

Table 26. Additional documentation resources for your system (continued)

Task	Document	Location
	<p>? > About.</p> <p>For information about updating drivers and firmware, see the Methods to download firmware and drivers section in this document.</p>	<p>Drivers</p>
<p>Managing your system</p>	<p>For information about systems management software offered by Dell, see the Dell OpenManage Systems Management Overview Guide.</p>	<p>OpenManage Documentation</p>
	<p>For information about setting up, using, and troubleshooting OpenManage, see the Dell OpenManage Server Administrator User's Guide.</p>	<p>OpenManage Documentation</p>
	<p>For information about installing, using, and troubleshooting Dell OpenManage Enterprise, see the Dell OpenManage Enterprise User's Guide.</p>	<p>OpenManage Documentation</p>
	<p>For information about partner programs enterprise systems management, see the OpenManage Connections Enterprise Systems Management documents.</p>	<p>OpenManage Connections and Integrations</p>
<p>Understanding event and error messages</p>	<p>For information about the event and error messages generated by the system firmware and agents that monitor system components, go to qrl.dell.com > Look Up > Error Code, type the error code, and then click Look it up.</p>	<p>PowerEdge Server Event and Error Messages</p>