

Dell EMC PowerEdge RAID Controller 10 User's Guide

PERC H345, H740P, H745, H745P MX, and H840

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: Overview.....	8
Supported operating systems.....	8
PERC card specifications.....	9
Management applications for PERC cards.....	10
Comprehensive Embedded Management.....	10
Dell OpenManage Storage Management.....	11
Chapter 2: Getting Started.....	12
PERC H345 front card.....	12
PERC H345 adapter.....	13
PERC H740P adapter.....	13
PERC H745 front card.....	14
PERC H745 adapter.....	15
PERC H745P MX card.....	15
PERC H840 adapter card.....	16
Chapter 3: Features.....	17
Controller features.....	17
1 MB I/O.....	17
Auto Configure RAID 0.....	17
Disk roaming.....	17
FastPath	18
PERC H745P MX enclosure support.....	18
Non-RAID disks.....	19
Physical disk power management.....	20
PERC H840 enclosure support.....	20
Profile Management.....	22
Secure firmware update.....	22
Snapdump.....	22
Virtual disk features.....	22
Virtual disk write cache policies.....	23
Virtual disk read cache policies.....	23
Virtual disk migration.....	24
Virtual disk initialization.....	24
Reconfiguration of virtual disks.....	25
Background operations.....	27
Hard drive features.....	27
Self-Encrypting Disks.....	27
Instant secure erase.....	28
4 KB sector disk drives.....	28
Fault tolerance.....	28
The SMART feature.....	28
Patrol Read.....	29
Physical disk failure detection.....	29

Using persistent hot spare slots.....	29
Physical disk hot swapping.....	30
Using replace member and revertible hot spares.....	30
Controller cache.....	30
Controller cache preservation.....	31
Battery Transparent Learn Cycle.....	31
Operating system device enumeration.....	31
Controller mode.....	32
RAID mode.....	32
Enhanced HBA mode.....	32
Non-RAID disk.....	33
Chapter 4: Deploying the PERC card.....	34
Removing the PERC H345 adapter.....	34
Installing the PERC H345 adapter.....	35
Removing the PERC H345 front card.....	36
Installing the PERC H345 front card.....	37
Removing the PERC H740P adapter.....	38
Installing the PERC H740P adapter.....	39
Removing the PERC H740P mini monolithic.....	39
Installing the PERC H740P mini monolithic.....	40
Removing the PERC H745 adapter.....	41
Installing the PERC H745 adapter.....	42
Removing the PERC H745 front card.....	43
Installing the PERC H745 front card.....	44
Removing the PERC H745P MX adapter card.....	45
Installing the PERC H745P MX adapter card.....	46
Removing the PERC H840 card.....	47
Installing the PERC H840 card.....	48
Part replacement of a PERC card in eHBA mode.....	48
Part replacement of a PERC card in RAID mode with another in eHBA mode.....	48
Chapter 5: Driver installation.....	50
Creating the device driver media.....	50
Downloading drivers from the Dell support website.....	50
Downloading drivers from the Dell systems service and diagnostic tools media	50
Windows driver installation.....	51
Installing the driver during a Windows Server 2012 R2 and newer installation.....	51
Installing the driver after Windows Server 2012 R2 and newer installation	51
Updating PERC 10 driver for existing Windows Server 2012 R2 and newer.....	52
Linux driver installation.....	52
Installing or updating the RPM driver package with KMOD support.....	53
Installing or updating the RPM driver package with KMP support.....	53
Loading the driver during operating system installation.....	53
Chapter 6: Firmware.....	55
Installing the firmware using DUP.....	55
Chapter 7: HII configuration utility.....	56

Entering the HII configuration utility.....	56
Exiting the HII configuration utility.....	56
Navigating to Dell PERC 10 configuration utility.....	56
HII Configuration utility dashboard view options.....	57
Configuration management.....	57
Auto Configure RAID 0.....	57
Creating virtual disks.....	58
Configuring virtual disk parameters.....	58
Creating profile based virtual disk.....	59
Viewing disk group properties.....	59
Deleting configurations.....	59
Controller management.....	59
Clearing controller events.....	59
Saving controller events.....	60
Saving debug log.....	60
Managing link speed.....	60
Restoring factory default settings.....	60
Manage controller mode.....	60
Changing profile using HII Configuration Utility.....	61
Advanced controller properties.....	61
Virtual disk management.....	64
Virtual disk numbering.....	64
Viewing virtual disk properties.....	64
Viewing physical disks associated with a virtual disk.....	65
Configuring virtual disk policies.....	65
Configuring Virtual Disks.....	66
Performing expand virtual disk operation.....	66
Performing consistency check.....	66
Physical disk management.....	67
Viewing physical disk properties.....	67
Cryptographic erase.....	68
Physical disk erase.....	68
Assigning global hot spare.....	69
Assigning dedicated hot spare.....	69
Convert to RAID capable.....	69
Convert to Non-RAID disk.....	70
Hardware components.....	70
Viewing battery properties.....	70
Viewing physical disks associated with an enclosure.....	71
Security key management in the HII configuration utility.....	71
Chapter 8: Security key and RAID management.....	72
Security key implementation.....	72
Local Key Management.....	72
Creating security key.....	72
Changing Security Settings.....	73
Disabling security key.....	73
Create secured virtual disk.....	73
Secure non-RAID disks.....	74
Secure preexisting virtual disk.....	74

Importing secure virtual disk.....	74
Import secured non-RAID disk.....	74
Dell EMC OpenManage Secure Enterprise Key Manager.....	75
Supported controllers for OpenManage Secure Enterprise Key Manager.....	75
Managing enterprise key manager mode.....	75
Disabling enterprise key manager mode.....	75
Managing virtual disks in enterprise key manager mode.....	75
Migration of drives from local key management to enterprise key management.....	76

Chapter 9: Troubleshooting..... 77

Single virtual disk performance or latency in hypervisor configurations.....	77
Adapter at baseport not responding error message.....	77
BIOS disabled error message.....	78
Configured disks removed or not accessible error message.....	78
Dirty cache data error message.....	78
Discovery error message.....	79
Drive Configuration Changes Error Message.....	79
Windows operating system installation errors	79
Firmware fault state error message.....	79
Extra enclosure error message.....	79
Foreign configuration found error message.....	80
Foreign configuration not found in HII error message.....	80
Degraded state of virtual disks.....	80
Memory errors.....	80
Preserved Cache State.....	81
Security key errors.....	81
Secured foreign import errors.....	81
Failure to select or configure non Self-Encrypting Disks non-SED.....	81
Failure to delete security key.....	81
Failure of Cryptographic Erase task on encryption-capable physical disks.....	81
General issues.....	82
PERC card has yellow bang in device manager.....	82
PERC card not seen in device manager.....	82
Physical disk issues.....	82
Physical disk in failed state.....	82
Unable to rebuild a fault tolerant virtual disk.....	82
Fatal error or data corruption reported.....	82
Physical disk displayed as blocked.....	82
Multiple disks are inaccessible.....	83
Rebuilding data for a failed physical disk.....	83
Virtual disk fails during rebuild using a global hot spare.....	83
Dedicated hot spare disk fails during rebuild.....	83
Redundant virtual disk fails during reconstruction.....	84
Virtual disk fails rebuild using a dedicated hot spare.....	84
Physical disk takes a long time to rebuild.....	84
SMART errors.....	84
Smart error detected on a physical disk in a redundant virtual disk.....	84
Smart error detected on a physical disk in a non-redundant virtual disk.....	85
Replace member errors.....	85
Source disk fails during replace member operation.....	85

Target disk fails during replace member operation.....	85
A member disk failure is reported in the virtual disk which undergoes replace member operation.....	85
Linux operating system errors.....	86
Virtual disk policy is assumed as write-through error message.....	86
Unable to register SCSI device error message.....	86
Drive indicator codes.....	86
HII error messages.....	87
Unhealthy Status of the Drivers.....	87
Rebuilding a drive during full initialization.....	87
System reports more drive slots than what is available.....	87
Chapter 10: Appendix RAID description.....	89
Summary of RAID levels.....	89
RAID 10 configuration.....	90
RAID terminology.....	91
Disk striping.....	91
Disk mirroring.....	91
Spanned RAID levels.....	92
Parity data.....	92
Chapter 11: Getting help.....	93
Recycling or End-of-Life service information.....	93
Contacting Dell.....	93
Locating the Express Service Code and Service Tag.....	93
Receiving automated support with SupportAssist	94
Chapter 12: Documentation resources.....	95

Overview

The PowerEdge RAID Controller (PERC) 10 series consist of the H345, H740P, H745, H745P MX, and H840 cards. The PERC 10 family of storage controller cards has the following characteristics:

- Complies with serial-attached SCSI (SAS) 3.0 providing up to 12 Gb/sec throughput.
- Supports Dell-qualified serial-attached SCSI (SAS) hard drives, SATA hard drives, and solid-state drives (SSDs).
- Offers RAID control capabilities including support for RAID levels 0, 1, 5, 6, 10, 50, and 60.
- Provides reliability, high performance, and fault-tolerant disk subsystem management.

NOTE: Mixing SAS and SATA drives within a virtual disk is not supported. Also, mixing hard drives and SSDs within a virtual disk is not supported.

NOTE: Mixing disks of different speed (7,200 rpm, 10,000 rpm, or 15,000 rpm) and bandwidth (3 Gbps, 6 Gbps or 12 Gbps) while maintaining the same drive type (SAS or SATA) and technology (HDD or SSD) is supported.

NOTE: RAID levels 5, 6, 50 and 60 are not supported on PERC H345.

NOTE: For the safety, regulatory, and ergonomic information associated with these devices, and for more information about the Integrated Dell Remote Access Controller (iDRAC) or Lifecycle Controller (LC) remote management, see your platform documentation.

Topics:

- [Supported operating systems](#)
- [PERC card specifications](#)
- [Management applications for PERC cards](#)

Supported operating systems

See [Dell EMC Enterprise operating systems support](#) for a list of supported operating systems by a specific server for the PERC 10 card.

NOTE: The PERC 10 driver for VMware ESXi is bundled with the VMware ISO image available from Dell. For more information, see www.dell.com/virtualizationsolutions.

NOTE: For the latest list of supported operating systems and driver installation instructions, see the operating system documentation at www.dell.com/operatingsystemmanuals. For specific operating system service pack requirements, see the Drivers and Downloads section at www.dell.com/manuals.

PERC card specifications

The table below lists and describes the different PERC cards that consist of the PERC 10 series and their specifications:

Table 1. PERC cards

Feature	PERC H345	PERC H740P	PERC H745	PERC H745P MX	PERC H840
RAID levels	0, 1, 10	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	N/A
Enclosures per port	Not applicable	Not applicable	Not applicable	Not applicable	4
Processor	Dell adapter SAS RAID-on-chip, 16-port with LSI 3416 chipset	Dell adapter SAS RAID-on-chip, 8-port with LSI 3508 chipset	Dell adapter SAS RAID-on-chip, 16-port with LSI 3516 chipset	Dell adapter SAS RAID-on-chip, 16-port with LSI 3516 chipset	Dell adapter SAS RAID-on-chip, 8-port with LSI 3508 chipset
Battery backup unit	No	Yes	Yes	Yes	Yes
eHBA	No	Yes	Yes	Yes	No
LKM security	No	Yes	Yes	Yes	Yes
Controller queue depth (RAID mode)	1536	4096	4096	3072	4096 with PD64 profile, 3072 with PD240 profile
Controller Queue Depth (eHBA)	N/A	5120	5120	5120	N/A
Enterprise key manager mode	No	Yes	Yes	No	No
Nonvolatile cache	No	Yes	Yes	Yes	Yes
Cache memory	N/A	8 GB DDR4 2133 Mhz cache	4 GB DDR4 2133 Mhz cache	8 GB DDR4 2133 Mhz cache	8 GB DDR4 2133 Mhz cache
Cache function	Write through and no read ahead.	Write back, write through, no read ahead, and read ahead.	Write back, write through, no read ahead, and read ahead.	Write back, write through, no read ahead, and read ahead.	Write back, write through, no read ahead, and read ahead.
Max no of VD's in RAID mode	32	64	64	240	240
Max no of disks in eHBA mode	Not applicable	240 (combination of virtual disks and Non-Raid disks).	240 (combination of virtual disks and Non-Raid disks)	240 (combination of virtual disks and Non-Raid disks).	Not applicable
Max no of VD's per disk group	16	16	16	16	16
Max no of disk groups	32	64	64	240	240
Hot swap devices supported	Yes	Yes	Yes	Yes	Yes
Hardware XOR engine	Not applicable	Yes	Yes	Yes	Yes
Online capacity expansion	Yes in operating system only	Yes	Yes	Yes	Yes
Dedicated and global hot spare	Yes	Yes	Yes	Yes	Yes
Autoconfigure	N/A	RAID 0	RAID 0	RAID 0	RAID 0

Table 1. PERC cards (continued)

Feature	PERC H345	PERC H740P	PERC H745	PERC H745P MX	PERC H840
Drives types	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS	6 Gbps SAS, and 12 Gbps SAS
VD strip size	64 kB	64 kB, 128 kB, 256 kB, 512 kB, 1 MB	64 kB, 128 kB, 256 kB, 512 kB, 1 MB	64 kB, 128 kB, 256 kB, 512 kB, 1 MB	64 kB, 128 kB, 256 kB, 512 kB, 1 MB
PCIe support	Gen 3	Gen 3	Gen 3	Gen 3	Gen 3
Maximum drive support	<ul style="list-style-type: none"> Without SAS Expander: 16 per controller With SAS Expander: Limited by platform offering 	<ul style="list-style-type: none"> Without SAS Expander: 8 drives per controller With SAS Expander: Limited by platform offering 	<ul style="list-style-type: none"> Without SAS Expander: 16 drives per controller With SAS Expander: Limited by platform offering 	<ul style="list-style-type: none"> Without SAS Expander: 16 drives per controller With SAS Expander: Limited by platform offering 	240

NOTE: Enhanced HBA mode is supported from firmware version 50.5.0-1750 and later.

NOTE: Enterprise key manager mode is supported from firmware version 50.5.1-2633 and later.

NOTE: 240 virtual disks can be created only within a PD240 profile, for more information see [Profile Management](#).

NOTE: For information on number of drives in a disk group per virtual disk, see [Summary of RAID levels](#).

Management applications for PERC cards

Dell OpenManage Storage Management applications 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. The management applications for all PERC cards include:

- Comprehensive Embedded Management
- Dell OpenManage Storage Management
- Human Interface Infrastructure (HII) Configuration Utility
- The PERC CLI

NOTE: The BIOS configuration utility <Ctrl> <R> is not supported on PERC 10 cards.

Comprehensive Embedded Management

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

Using CEM enables you to perform the following:


- Monitor devices with and without an operating system installed on the system.
- Provide a specific location to access monitored data of the storage devices and network cards.
- Allows controller configuration for all PERC 10 cards.

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.

NOTE: It is not recommended to create more than 8 VDs simultaneously with CEM.

Dell OpenManage Storage Management

Dell OpenManage Storage Management is a storage management application for Dell systems that provides enhanced features for configuring locally attached RAID disk storage. The Dell OpenManage storage management application enables you to perform controller and enclosure functions for all supported RAID controllers and enclosures from a single graphical or command-line interface (CLI). The user interface (UI) is wizard-driven with features for novice and advanced users, and detailed online help. 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 select operating systems, allows you to perform RAID management tasks either directly from the console or through scripting.

 **NOTE:** For more information, see the Dell OpenManage Storage Management User's Guide at www.dell.com/openmanagemanuals.

Getting Started

The Dell PowerEdge RAID Controller (PERC) 10 series of cards consist of the H345, H740P, H745, H745P MX, and H840 cards.

Topics:

- PERC H345 front card
- PERC H345 adapter
- PERC H740P adapter
- PERC H745 front card
- PERC H745 adapter
- PERC H745P MX card
- PERC H840 adapter card

PERC H345 front card




Figure 1. Features of PERC H345 front card

- | | |
|-----------------------------------|-----------------------------------|
| 1. PCIe cable connector | 2. Heat sink |
| 3. SAS/SATA backplane connector B | 4. SAS/SATA backplane connector A |
| 5. Power connector | |

PERC H345 adapter




Figure 2. Features of PERC H345 adapter

- 1. Heat sink
- 2. SAS/SATA backplane connector B
- 3. SAS/SATA backplane connector A
- 4. PCIe connector

PERC H740P adapter

The PERC H740P is the performance RAID solution card consisting of 8 GB non-volatile cache and is available in the adapter (low profile and full height) and mini monolithic form factors for internal storage.




Figure 3. Features of PERC H740P adapter

- 1. Heat sink
- 2. Battery
- 3. Battery cable
- 4. Battery-cable connector
- 5. Port B
- 6. Port A




Figure 4. Features of PERC H740P mini monolithic card

1. Heat sink
2. Battery
3. Battery cable

PERC H745 front card




Figure 5. Features of PERC H745 front card

- | | |
|--|---|
| <ol style="list-style-type: none"> 1. PCIe cable connector 3. Heat sink 5. SAS/SATA backplane connector B 7. Power connector | <ol style="list-style-type: none"> 2. Battery 4. Battery cable connector 6. SAS/SATA backplane connector A |
|--|---|

PERC H745 adapter




Figure 6. Features of PERC H745 adapter

- | | |
|-----------------------------------|-----------------------------------|
| 1. Heat sink | 2. Battery |
| 3. Battery cable connector | 4. SAS/SATA backplane connector B |
| 5. SAS/SATA backplane connector A | 6. PCIe connector |

PERC H745P MX card

The PERC H745P MX is the RAID solution card for the MX7000 chassis. 8 GB of non-volatile cache manages the drives from its local enclosure and from the MX5016s storage enclosure.




Figure 7. Features of PERC H745P MX card

- | | |
|----------------------------|------------------------|
| 1. Mezzanine connector | 2. Release lever |
| 3. Heat sink | 4. SAS-cable connector |
| 5. Battery-cable connector | 6. Battery |

PERC H840 adapter card

The PERC H840 is similar to the H740P solution, except that it supports external storage. The PERC H840 is only available in the Adapter (low profile and full height) form factor.




Figure 8. Features of PERC H840 adapter card

- 1. Heat sink
- 2. Battery
- 3. Battery cable
- 4. Battery-cable connector
- 5. Port B/1
- 6. Port A/O

Features

Topics:

- Controller features
- Virtual disk features
- Hard drive features
- Fault tolerance
- Operating system device enumeration
- Controller mode
- Non-RAID disk

Controller features

This section lists the following controller features supported on PowerEdge RAID Controller (PERC) 10 cards in detail:

- 1 MB I/O
- Auto Configure RAID 0
- Disk roaming
- FastPath
- H745P MX enclosure support
- Non-RAID disk
- Physical disk power management
- PERC H840 enclosure support
- Profile Management
- Secure firmware update
- Snapdump

1 MB I/O

PERC 10 controllers support a 1 MB I/O feature; if the capacity of I/O frame is greater than 1 MB, the I/O frame is broken into smaller chunks.

Auto Configure RAID 0

Auto Configure RAID 0 feature creates a single drive RAID 0 on each physical disk that is in the ready state. For more information, see [Auto Configure RAID 0](#).

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 places them in the virtual disks that are part of the disk group.

 **CAUTION:** It is recommended that you perform disk roaming when the system is turned off.

 **CAUTION:** Do not attempt disk roaming during RAID level migration (RLM) or online capacity expansion (OCE). This causes loss of the virtual disk.

Using disk roaming

Perform the following steps to use disk roaming:

1. Turn off the power to the system, physical disks, enclosures, and system components.
2. Disconnect power cables from the system.
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. Turn on the system.

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

FastPath

FastPath is a feature that improves application performance by delivering high I/O per second (IOPs) for solid-state drives (SSDs). The Dell PERC 10 series supports FastPath.

To enable FastPath on a virtual disk, the cache policies of the RAID controller must be set to write-through and not read ahead. This enables FastPath to use the proper data path through the controller based on command (read/write), I/O size, and RAID type. For optimal solid-state drive performance, create virtual disks with strip size of 64 KB.

PERC H745P MX enclosure support

The H745P MX is a RAID controller for the PowerEdge MX740c and MX840c compute sleds. It is used as an interface in a multipath topology with the MX5016s storage sled for the MX7000 modular system. The following illustration describes the multipath topology for the MX7000 modular system:

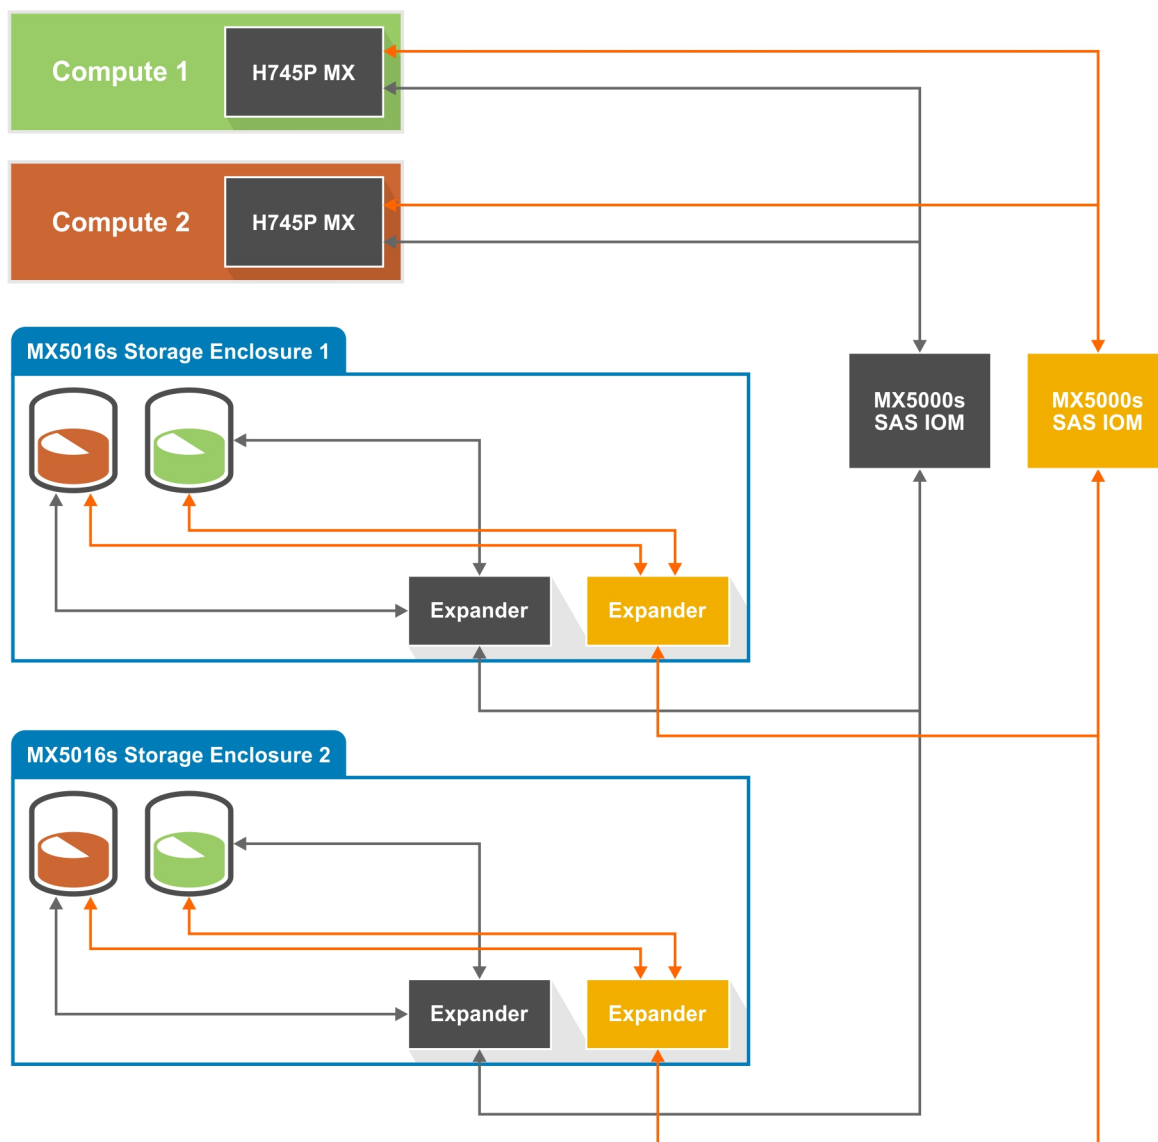


Figure 9. Multipath topology for the MX7000 modular system

In this illustration, the H745P MX resides in the compute sled and functions as an interface to the SAS topology of the MX7000. The H745P MX is used to configure drives that are located in the MX5016s storage sled into RAID volumes. The H745P MX connects to a pair of MX5000s SAS I/O modules (IOMs) that manage connections between the compute and storage sleds in the MX7000. For more information about configuring drive assignments, see the www.dell.com/openmanagemanuals > OpenManage Enterprise Modular.

NOTE: The I/O modules cannot be configured independently.

The H745P MX is connected to the SAS IOM, which in turn connects to one of the SAS expanders in the storage sled. The expander in the storage sled is then connected to all the physical drives in the storage sled, providing one path between the H745P MX and a physical drive in the storage sled. With connections to both SAS IOMs, the H745P MX has two paths to each drive in a storage sled allowing for a multi-path configuration. If either a SAS IOM or a SAS expander in the storage sled is removed, then a path to the physical drive is removed from the SAS topology.

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.

Physical disk power management

Physical disk power management is a power-saving feature of PERC 10 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. The physical disk power management feature is disabled by default. You can enable the feature in the Dell Open Manage Storage Management application or in the Human Interface Infrastructure (HII) configuration utility. For more information on HII configuration and physical disk power management, see [Enabling physical disk power management](#). For more information on using the Dell Open Manage Storage Management application, see the Dell OpenManage documentation at www.dell.com/openmanagemanuals.

PERC H840 enclosure support

The H840 card is an external card. The H840 card supports:

- Up to eight enclosures using single path configuration.
- Up to four enclosure using multipath configuration.

To set up a configuration with redundant paths, both ports on a controller must be cabled to the ports of a single enclosure.

NOTE: The PERC H840 card supports redundant paths when used with Dell PowerVault MD14XX series disk storage enclosures.

NOTE: The Enhanced HBA mode is not supported on PERC H840.




Figure 10. PERC H840 card ports—Four enclosures with multipath configuration

To set up a multipath enclosure on the PERC H840 card:

1. To connect a single enclosure, cable the ports so that the enclosure is connected to both the ports on the controller card.
2. To connect multiple enclosures, connect two cables from the controller to the first enclosure as in case in point no 1. Next, connect cables from the first enclosure to the second enclosure. From here, continue to connect two cables from the last enclosure in the chain to the next enclosure you wish to include until you have up to 4 enclosures connected.

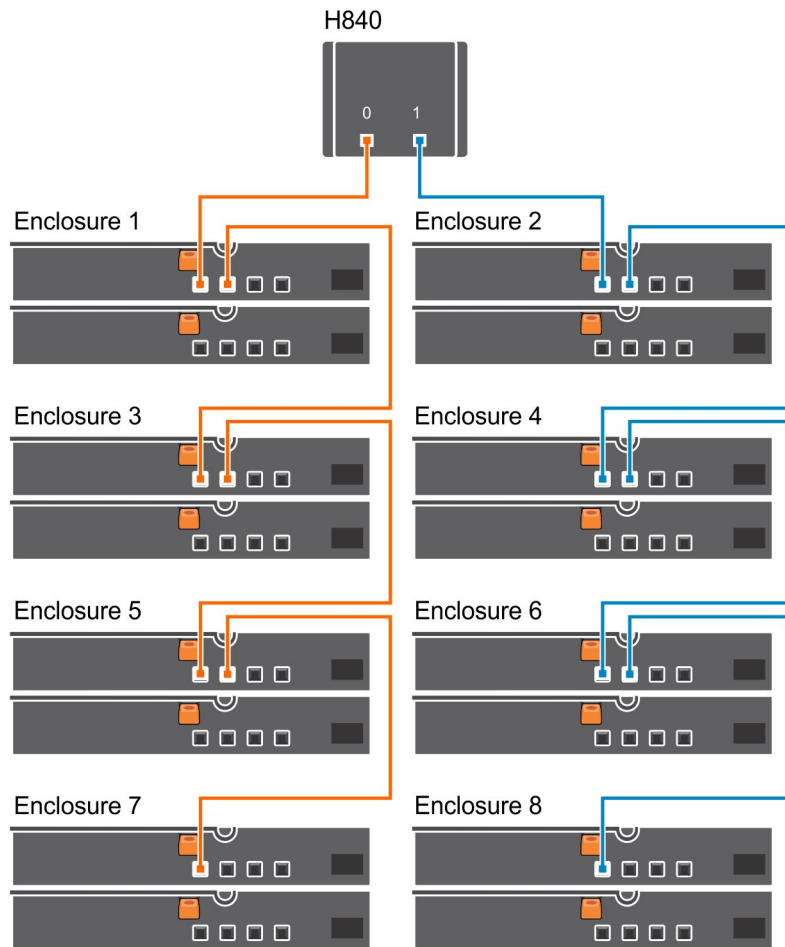


Figure 11. PERC H840 card ports—Eight enclosures with single path configuration

To set up a single path enclosure on the PERC H840 card:

1. To connect a single enclosure, cable the ports so that the enclosure is connected to both the ports on the controller card.
2. To connect multiple enclosures, single cables connect the first enclosure to the second enclosure, then the single cable connects the second enclosure to the third enclosure.

NOTE: You must use a PD240 profile to support 4 enclosures per port. For more information, see [Profile Management](#).

Redundant path support for PERC H840

The PERC H840 adapter can detect and use redundant paths to disks contained in enclosures. This provides the ability to connect two SAS cables between a controller and an enclosure for path redundancy. The controller can tolerate the failure of a cable or enclosure management module (EMM) by utilizing the remaining path. When redundant paths exist, the controller automatically balances I/O loads through both paths to each disk. Load balancing increases throughput to virtual disks in storage enclosures and is automatically turned on when redundant paths are detected. To set up your hardware to support redundant paths, see [Setting up redundant path support on the PERC H840 adapter](#).

NOTE: This is applicable for the PERC H840 only.

NOTE: This support for redundant paths refers to path redundancy only and not to controller redundancy.

Setting up redundant path support on the PERC H840 adapter

The PERC H840 card can detect and use redundant paths to disks contained in enclosures. With redundant paths to the same device, if one path fails, another path can be used to communicate between the controller and the device.

To connect single enclosure, cable the ports so that the enclosure is connected to both the ports on the controller card. See **Figure 6. PERC H840 Card ports - Four enclosures with multipath configuration** in PERC H840 enclosure support.

Profile Management

PERC 10 introduces the concept of profiles. A profile is user selectable. It defines controller queue depth and the maximum number of physical and virtual disks. There are two profiles: PD64 and PD240. The H740 supports only the PD64 profile. The H745P MX supports only the PD240 profile. The H840 supports both PD64 and PD240. PD240 is the default profile on H840. For more information see [Changing profile using HII Configuration Utility](#).

NOTE: When an H740 or H840 controller is upgraded to firmware version 50.3.0-1022 from an older firmware version, the profile on the controller is set to PD64 by default.

Table 2. Supported Profiles on PERC 10

Feature	PD64	PD240
Profile ID	10	12
Maximum physical disk supported	64	240
Maximum virtual disk supported	64	240
Controller queue depth	4096	3072
Maximum enclosure supported	<ul style="list-style-type: none">One enclosure using multipath configurationTwo enclosure using single path configuration	<ul style="list-style-type: none">Four enclosure using multipath configurationEight enclosure using single path configuration
NOTE: In the MX chassis, H745P MX supports seven enclosures using a multipath configuration.		
NOTE: PD64 is supported on both internal and external controllers, whereas PD240 is supported only on the H840 and H745P MX.		

Secure firmware update

This feature provides a cryptographic method of updating the firmware using an RSA encryption-decryption algorithm.

Only Dell-certified firmware is supported on your PERC controller.

Snapdump

The Snapdump feature provides the Dell support team with the debug information which 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. This means 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 reboots before its deleted.

To generate a snapdump, change the snapdump, delete a snapdump, and to download a stored snapdump settings, see *Dell EMC PowerEdge RAID Controller CLI Reference Guide* at www.dell.com/storagecontrollermanuals.

Virtual disk features



This section lists the following virtual disk features supported on PERC 10 cards in detail:


- Single virtual disk performance or latency in hypervisor configurations
- Virtual disk write cache policy
- Virtual disk read cache policy
- Virtual disk migration
- Virtual disk Initialization
- Reconfiguration of virtual disk
- Background operations


Virtual disk write cache policies

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

Table 3. Write cache policies

Feature	Description
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. The controller then writes the cached data to the storage device in the background.  NOTE: The default cache setting for virtual disks is Write-back caching. Write-back caching is also supported for single drive RAID 0 virtual disks.
Write-through	The controller sends a data transfer completion signal to the host system when the disk subsystem has received all the data in a transaction.  NOTE: Certain data patterns and configurations perform better with a write-through cache policy.

 **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.

 **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 employed

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

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

 **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.

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

Virtual disk read cache policies

The read policy of a virtual disk determines how the controller handles reads to that virtual disk.

Table 4. Read policies

Feature	Description
Read ahead	Allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data is required soon. This speeds up reads for sequential data, but there is slight improvement when accessing random data.
No read ahead	Disables the read ahead capability.

Table 4. Read policies (continued)


Feature	Description
Adaptive read ahead	Adaptive read ahead is no longer supported. Selecting adaptive read ahead is equivalent to selecting the read ahead option.

Virtual disk migration


The PERC 10 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 flags the physical disk as foreign, and generates an alert indicating that a foreign disk was detected.


Disk migration pointers:


- Supports migration of virtual disks from PERC H330, H730, H730P, and H830 to the PERC 10 series except for PERC H345.
- Supports migration of volumes created within the PERC 10 series.
- Does not support migration from the PERC 10 series to PERC H330, H730, H730P, H830, H310, H710, H710P, H810.
- Does not support migration from PERC H310, H710, H710P, and H810 to the PERC10 series.

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

 **NOTE:** Importing non-RAID drives and uneven span RAID 10 virtual disks from PERC 9 to PERC 10 is not supported.

 **NOTE:** Disks cannot be migrated to older generations of PERC cards.

 **NOTE:** Importing secured virtual disks is supported as long as the appropriate local key management (LKM) is supplied or configured.

 **NOTE:** RAID 5, 6, 50, and 60 virtual disks cannot be imported in eHBA mode.


 **NOTE:** Virtual disk migration from PERC H740P, H745P, H745, H840 to H345 is not supported.

 **CAUTION:** Do not attempt disk migration during RLM or online capacity expansion (OCE). This causes loss of the virtual disk.

Virtual disk initialization

PERC 10 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

Performing a full initialization on a virtual disk overwrites all blocks and destroys 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.

You can start a full initialization on a virtual disk by using the Slow Initialize option in the Dell OpenManage storage management application. For more information on using the HII Configuration Utility to perform a full initialization, see [Configuring virtual disk parameters](#).

 **NOTE:** If the system reboots during a full initialization, the operation aborts and a BGI begins on the virtual disk.

Fast initialization

A fast initialization on a virtual disk overwrites the first and last 8 MB 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 a longer time to complete. To perform a fast initialization using the HII Configuration Utility, see [Configuring virtual disk parameters](#).

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 create a VD, 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, then the I/O operation fails and the operating system reports an I/O error.

Reconfiguration of virtual disks

An online virtual disk can be reconfigured in ways that expands its capacity and changes its RAID level.

NOTE: Spanned virtual disks such as RAID 50 and 60 cannot be reconfigured.

NOTE: Reconfiguring virtual disks typically impacts disk performance until the reconfiguration operation is complete.

Online Capacity Expansion (OCE) can be done in two ways:

1. If there is a single virtual disk in a disk group and free space is available, the capacity of a virtual disk can be expanded within that free space. If multiple virtual disks exist within a common disk group, the capacities of those virtual disks cannot be expanded.

NOTE: Online capacity expansion is allowed on a disk group with a single virtual disk that begins at the start of the physical disk. It is not allowed when there is a free space at the beginning of a disk.

2. Free space is also available when the physical disks of a disk group are replaced by larger disks using the replace member feature. The capacity of a virtual disk can also be expanded by performing an OCE operation to add more physical disks.

RAID level migration (RLM) refers to changing a virtual disk's RAID level. Both RLM and OCE can be done simultaneously so that a virtual disk can simultaneously have its RAID level that is changed and its capacity increased. When an RLM or an OCE operation is complete, a reboot is not required.

CAUTION: Do not attempt disk migration during RLM or OCE operations. This causes loss of the virtual disk.

NOTE: If an RLM or an OCE operation is in progress, then an automatic drive rebuild or copyback operation will not start until the operation is complete.

NOTE: If the controller already contains the maximum number of virtual disks, you cannot perform a RAID level migration or capacity expansion on any virtual disk.


NOTE: The controller changes the write cache policy of all virtual disks to write-through until the RLM or OCE operation is complete.

NOTE: You cannot initiate an OCE or an RLM on any virtual disk on the controller where a virtual disk with an ID of 239 exists.

See the following table for a list of RLM or OCE options: The source RAID level column indicates the virtual disk RAID level before the RLM or OCE operation and the target RAID level column indicates the RAID level after the RLM or OCE operation.

Table 5. RAID level migration

Source RAID Level	Target RAID Level	Number of Physical Disks (Beginning)	Number of Physical Disks (End)	Capacity Expansion Possible	Description
RAID 0	RAID 0	1 or more	2 or more	Yes	Increases capacity by adding disks.
RAID 0	RAID 1	1	2	Yes	Converts a non-redundant virtual disk into a mirrored virtual disk by adding one disk.
RAID 0	RAID 5	1 or more	3 or more	Yes	Adds distributed parity redundancy; at least one disk needs to be added.
RAID 0	RAID 6	1 or more	4 or more	Yes	Adds dual distributed parity redundancy; at least two disks need to be added.
RAID 1	RAID 0	2	2 or more	Yes	Removes redundancy while increasing capacity.
RAID 1	RAID 5	2	3 or more	Yes	Maintains redundancy while adding capacity.
RAID 1	RAID 6	2	4 or more	Yes	Adds dual distributed parity redundancy and adds capacity.
RAID 5	RAID 0	3 or more	2 or more	Yes	Converts to a non-redundant virtual disk and reclaims disk space that is used for distributed parity data; one disk can be removed.
RAID 5	RAID 5	3 or more	4 or more	Yes	Increases capacity by adding disks.
RAID 5	RAID 6	3 or more	4 or more	Yes	Adds dual distributed parity redundancy; at least one disk needs to be added.
RAID 6	RAID 0	4 or more	2 or more	Yes	Converts to a non-redundant virtual disk and reclaims disk space that is used for distributed parity data; two disks can be removed.
RAID 6	RAID 5	4 or more	3 or more	Yes	Removes one set of parity data and reclaims disk space used for it; one disk can be removed.
RAID 6	RAID 6	4 or more	5 or more	Yes	Increases capacity by adding disks.
RAID 10	RAID 10	4 or more	6 or more	Yes	Increases capacity by adding disks; an even number of disks need to be added.

 **NOTE:** You cannot perform a RAID level migration and expansion on RAID levels 50 and 60.

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 Dell OpenManage storage management application. Any change to the BGI rate does not take effect until the next BGI is executed.

NOTE:

- On PERC H345, background operations will not run until the operating system boots.
- You cannot disable BGI permanently. 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 Dell OpenManage storage management application. You can schedule a CC to run on virtual disks using the Dell OpenManage storage management application. To start a CC using the HII Configuration Utility, see [Performing consistency check](#).

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.

Hard drive features

This section lists the following hard drive features supported on PERC 10 cards in detail:

- Self-encrypting disks (SED)
- Instant secure erase (ISE)
- 4 KB sector disk drives

Self-Encrypting Disks

The Dell PERC 10 series of cards support self-encrypting disks (SED) 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 under local key management (LKM). The security key is used by the controller to lock and unlock access to encryption-capable physical disks. In order to take advantage of this feature, you must:

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

For more information, see the [Security key and RAID management](#) section.

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

NOTE: PERC supports Trusted Computing Group Enterprise (TCG) Security Subsystem Classes (SSC) SAS/SATA SED drives.

Instant secure erase

Instant secure erase (ISE) drives use the same encryption technology as SED drives 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.

NOTE: ISE drives do not provide protection against theft.

4 KB sector disk drives

PERC 10 controllers support 4 KB sector disk drives, which enables you to efficiently use the storage space.

Before installing Windows on 4 KB sector disk drives, refer to [Windows operating system installation errors](#).

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

NOTE: 4 K is only supported in UEFI mode and not legacy BIOS.

NOTE: 4 K devices do not appear in the boot select option.

Fault tolerance

The PERC 10 series supports the following:

- Self-monitoring and reporting technology (SMART)
- Patrol read
- Physical disk failure detection
- Physical disk rebuild using hot spares
- Controller cache preservation
- Battery and non-volatile cache backup of controller cache to protect data
- Detection of batteries with low charge after boot up

The next sections describe some 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 that are predicted physical disk failures, such as a bearing failure, a broken read/write head, and changes in spin-up rate. In addition, there are factors related to read/write surface failure, such as seek error rate and excessive bad sectors.

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

Automatic Replace Member with predicted failure

A replace member operation can occur when there is a SMART predictive failure reporting on a physical disk in a virtual disk. The automatic replace member 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 replace member. This prevents the array from reaching degraded state.

If an automatic replace member 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 replace member operation successfully completes.

NOTE: To enable automatic replace member, use the Dell OpenManage storage management application.

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.

The following is an overview of patrol read behavior:

- Patrol read runs on all disks on the controller that are configured as part of a virtual disk, including hot spares.
- Patrol read does not run on physical disks that are not part of a virtual disk or 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
 - CC
 - RLM or OCE

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

For more information about patrol read, see the Dell OpenManage documentation at www.dell.com/openmanagemanuals.

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 occur with 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

NOTE: The persistent hot spare slot feature is disabled by default.

The PERC 10 series can be configured so that the system backplane or storage enclosure disk slots are dedicated as hot spare slots. This feature can be enabled using the Dell OpenManage storage management application.

Once enabled, any slots with hot spares configured automatically become persistent hot spare slots. 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 one 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 on persistent hot spares, see the Dell OpenManage documentation at www.dell.com/openmanagemanuals.

Configured slot behavior

This feature is similar to persistent hot spare slot behavior. If a redundant VD is configured to the system and if a drive is replaced, the configured slot will automatically rebuild or copyback on the inserted drive regardless of the data on the drive. This operation will overwrite the data on the drive.

Table 6. Drive state/operation

Drive state/operation	Unconfigured slot	Slot configured in VD
Insert unconfigured drive into the system	Ready	Rebuild or copyback start
Insert configured drive into the system	Foreign	Rebuild or copyback start <ul style="list-style-type: none">• Original drive data lost

Table 6. Drive state/operation (continued)

Drive state/operation	Unconfigured slot	Slot configured in VD
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 start• Original drive data lost
Insert locked drive into the system (non-unlockable)	Foreign locked	Foreign locked

Physical disk hot swapping

NOTE: To check if the backplane supports hot swapping, see your system owner's manual.

Hot swapping is the manual replacement of a disk while the PERC 10 series cards are online and performing their normal functions. The following requirements must be met before hot swapping a physical disk:

- The system backplane or enclosure must support hot swapping for the PERC 10 series cards.
- The replacement disk 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.

Using replace member and revertible hot spares

The replace member functionality allows 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 replace member operation typically causes a temporary impact to disk performance. Once the operation completes, performance returns to normal.

Controller cache

The H740P, H745, H745P MX, and H840 cards contain the local DRAM on the controllers. This DRAM can cache I/O operations per second for Write Back, Read Ahead virtual disks to improve the performance.

NOTE: Virtual disks consisting of SSDs may not see a difference in performance using controller cache and may benefit by [Fastpath](#).

I/O workload that is slow to HDDs, such as random 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 amount of data in the cache.

The following operations require a complete cache flush:

CAUTION: Reduce or stop the running workload before performing any of these operations.

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

NOTE: The iDRAC or OpenManage periodically scans for the foreign configurations when the foreign disks are present, these actions may degrade the performance periodically. If a foreign disk is present, it is recommended to 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 in the event of a system power outage or improper system shutdown. The PERC 10 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, a small amount of power from the battery is used to transfer the cache data to non-volatile flash storage where it remains until power is restored and the system is booted. If the cache preservation process is interrupted by power-on, the controller may request an extra reset during the boot to complete the process. The system displays a message during boot as `Dell PERC at Bus <X> Dev <Y> has requested a system reset. System will reboot in 5 seconds.`

Recovering cache data


If a system power loss or improper system shutdown has occurred:

1. Restore the system power.
2. Boot the system.
3. When preserved cache exists on the controller, an error message is shown. For more information on how to recover cache, see [Preserved Cache State](#).

Battery Transparent Learn Cycle

A transparent learn cycle is a periodic operation that calculates the charge that is remaining in the battery 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 the Transparent Learn Cycle (TLC) on the battery to calibrate and gauge its charge capacity once every 90 days. The operation can be performed manually if required.

 **NOTE:** Virtual disks stay in write-back mode, if enabled, during transparent learn cycle. When the TLC completes, the controller sets the next TLC to +90 days.

Transparent Learn Cycle completion time

The time frame for completion of a learn cycle is a function of the battery 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 mid cycle, it begins at a new cycle.

Conditions for replacing the battery


The PERC battery is marked failed when the state or health of the battery is declared bad. If the battery is declared failed, then all the virtual disks in write-back mode transitions to write-through mode, and the firmware runs learn cycles in subsequent reboots until the battery is replaced. On replacing the battery, virtual disk transitions to write-back mode.

Operating system device enumeration

Virtual disks and non-RAID disks are presented to the operating system as SCSI devices. The operating system enumerates these devices based on the SCSI target device in a particular logical order.

The order for PERC H345 is:

1. Non-RAID disks are enumerated first based on the physical device ID.
2. Virtual disks are enumerated second based on virtual disk target ID.

 **NOTE:** H345 non-RAID disks may not appear in slot order.

The order for H740P, H745, H745P MX, H840 is:

1. Non-RAID disks are enumerated first based on slot ID.
2. Virtual disks are enumerated second based on virtual disk target ID.

NOTE: If a virtual disk target ID occupies the slot ID, a non-RAID disk may not display in the slot order.

NOTE: OS 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 reboot the system for the final device enumeration after creating any virtual disks or non-RAID disks.

HII and PERC CLI can be used to identify non-RAID disk slot ID and device ID or virtual disk target ID. See the [Virtual disk management](#) and [Physical disk management](#) sections for details on HII. For additional PERC CLI details, see the [Dell EMC PowerEdge RAID Controller Command Line Interface Reference Guide](#).

Controller mode

You can configure the controllers in two modes.

- RAID mode
- Enhanced HBA mode

RAID mode

RAID mode is the default mode for PERC 10 series controllers. In this mode, all RAID levels are accessible to the user for virtual disk creation. Only online virtual disks are presented to the host, and all unconfigured or foreign disks are hidden from the host. All controller models (H345, H740P, H745, H745P MX, and H840) support RAID mode.

NOTE: In H345 controller, non-RAID disks can be created and assigned as boot devices in RAID mode. The H345 Controller only supports RAID levels 0,1,10. See [PERC Card Specifications](#) for more details.

Enhanced HBA mode

Enhanced HBA mode or eHBA mode is an alternative mode for PERC 10 series controllers.

NOTE: Older versions of OpenManage Storage Management and Comprehensive Embedded Management do not support management of the controller in the eHBA mode.

NOTE: In the eHBA mode, non-RAID disks are enumerated first in ascending order, while RAID volumes are enumerated in descending order. For more information, see the applicable systems management documentation.

NOTE: The only controllers that support eHBA mode are H740P, H745, and H745P MX. eHBA mode is not supported on H345 and H840.

The major differences and similarities between RAID mode and eHBA mode are listed in the following table:


Controller operating mode	Enhanced HBA mode (H740, H745, H745P MX)	RAID mode (H740, H745)	RAID mode (H745P MX, H840)	RAID mode (H345)
Supported RAID levels	0,1,10	0,1,5,6,10,50,60	0,1,5,6,10,50,60	0, 1, 10
Present non-RAID disks to the host. (see Non-RAID disk)	Yes	No	No	Yes
Automatically converts all unconfigured disks to non-RAID disks during system boot. controller	Yes	No	No	No

Controller operating mode	Enhanced HBA mode (H740, H745, H745P MX)	RAID mode (H740, H745)	RAID mode (H745P MX, H840)	RAID mode (H345)
reset, and when unconfigured disks are hot-inserted.				
Write back-caching	Yes (VDs only)	Yes	Yes	No
Key management support (see Security key and RAID management)	LKM	SEKM and LKM	LKM	No
Patrol read	Yes (VDs only)	Yes	Yes	Yes (VDs only)

Non-RAID disk


A non-RAID disk is a single disk that is presented and accessible to the host for I/O. It is not a RAID volume. The only supported cache policy for non-RAID disks is write-through.

H740P, H745, and H745P MX non-RAID disks are only configurable in the eHBA mode. They are not configurable in RAID mode. For more details on eHBA mode, see [Enhanced HBA mode](#) and for instructions on how to configure a non-RAID disk, see [Convert to non-RAID disk](#). H345 does support non-RAID drives in RAID mode.

 **NOTE:** H840 does not support non-RAID drives.

Deploying the PERC card


This section provides a set of high-level installation and removal instructions for the Dell PowerEdge RAID Controllers (PERC) 10 series.


 **NOTE:** For detailed information on cabling the PERC 10 cards, see the system documentation at <https://www.dell.com/poweredgemanuals>.

Topics:

- [Removing the PERC H345 adapter](#)
- [Installing the PERC H345 adapter](#)
- [Removing the PERC H345 front card](#)
- [Installing the PERC H345 front card](#)
- [Removing the PERC H740P adapter](#)
- [Installing the PERC H740P adapter](#)
- [Removing the PERC H740P mini monolithic](#)
- [Installing the PERC H740P mini monolithic](#)
- [Removing the PERC H745 adapter](#)
- [Installing the PERC H745 adapter](#)
- [Removing the PERC H745 front card](#)
- [Installing the PERC H745 front card](#)
- [Removing the PERC H745P MX adapter card](#)
- [Installing the PERC H745P MX adapter card](#)
- [Removing the PERC H840 card](#)
- [Installing the PERC H840 card](#)
- [Part replacement of a PERC card in eHBA mode](#)
- [Part replacement of a PERC card in RAID mode with another in eHBA mode](#)

Removing the PERC H345 adapter

 **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.

1. Turn 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 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. Disconnect any SAS cables connected to the card:
 - a. Press down and hold the metal tab on the SAS cable connector.
 - b. Pull the SAS cable out of the connector.
6. Replace the storage controller and reconnect the SAS cable before placing them in the riser. For more information on installing the card, see [Installing the PERC H345 adapter](#).

7. Reinstall the riser on the system board and fasten the riser.
8. Close the system.
9. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.




Figure 12. Removing and installing the PERC H345 adapter

- | | |
|---------------------------------------|-----------------------------------|
| 1. Heat sink | 2. SAS/SATA backplane connector B |
| 3. SAS/SATA backplane connector A | 4. PCIe connector |
| 5. Card connector on the system board | |

Installing the PERC H345 adapter

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.

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 system board.

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

4. Press the card-edge down until the card is fully seated.
5. Connect the SAS data 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.

6. Route the SAS data cable through the channel on the inner side of the chassis to the backplane.

7. Attach the connector labeled SAS A to connector SAS A on the backplane, and attach the connector labeled SAS B to connector SAS B on the backplane.
8. Close the system.
9. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.

Removing the PERC H345 front card

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.

1. Turn 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 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.
 - a. Systems with an H345 front controller in the upside-down orientation will need to have both the backplane and controller removed at the same time due to the small amount of clearance available. To do so, uninstall all drives from the backplane, disconnect all cables to the PERC and backplane, and then lift the backplane and PERC out of the system carefully.
5. Disconnect any SAS cables connected to the card:
 - a. Press down and hold the metal tab on the SAS cable connector.
 - b. Pull the SAS cable out of the connector.
6. Remove the PERC controller from the controller carrier.
7. Insert the replacement controller into the carrier and secure it with the appropriate screws.
8. Take the replacement storage controller and reconnect the SAS cable before reconnecting it to the backplane.
 - a. For systems configured with a H345 front controller in an upside down orientation, reattach the PERC controller to the backplane first before reinstalling the backplane into the system. For more information on installing the card, see [Installing the PERC H345 front controllers](#).
9. Close the system.
10. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.




Figure 13. Removing and installing the PERC H345 front card

- | | |
|-----------------------------------|-----------------------------------|
| 1. PCIe cable connector | 2. Heat sink |
| 3. SAS/SATA backplane connector B | 4. SAS/SATA backplane connector A |
| 5. Power connector | |

Installing the PERC H345 front card

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.

1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.
2. Open the system.
3. Connect the PERC card to the carrier and ensure the screws are properly fastened in place.

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

4. Align the carrier with the guide pins until the controller is securely seated.
5. Slide the card until it is fully seated in the connector. Tighten the screws on the carrier that connect to the chassis to secure the carrier.
6. Connect the SAS data 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.

7. Close the system.
8. Reconnect the system to its electrical outlet and turn on the system and any attached peripherals.

Removing the PERC H740P adapter

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.

1. Turn 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. Lift the card to remove it from the connector on the system board.
5. Disconnect the SAS cables connected to the card:
 - a. Press down and hold the metal tab on the SAS cable connector.
 - b. Pull the SAS cable out of the connector.
6. Replace the storage controller card and connect the cable. For more information on installing the card, see [Installing the PERC H740P adapter](#).
7. Close the system.
8. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.




Figure 14. Removing and installing the PERC H740P adapter

- | | |
|---------------------------------------|-----------------------------|
| 1. Heat sink | 2. Battery |
| 3. Battery cable | 4. SAS cable connectors (2) |
| 5. SAS cables (2) | 6. PERC card connector |
| 7. Card connector on the system board | |

Installing the PERC H740P adapter

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.

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 system board.

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

4. Press the card-edge down until the card is fully seated.
5. Connect the SAS data 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.

6. Route the SAS data cable through the channel on the inner side of the chassis to the backplane.
7. Attach the connector labeled "SAS A" to connector SAS A on the backplane, and attach the connector labeled "SAS B" to connector SAS B on the backplane.
8. Close the system.
9. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.

Removing the PERC H740P mini monolithic

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.

1. Turn 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. Using a #2 Philips screw driver, loosen the screws that secure the controller cable to the card connector on the system board.
5. Lift the controller cable away from the card.
6. Lift the connector end of the card and angle it to disengage the card from the card holder on the system board.
7. Lift the card out of the system.
8. Close the system.
9. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.




Figure 15. Removing and installing the PERC H740P mini monolithic card

- | | |
|---------------------------------------|------------------------------------|
| 1. Cable | 2. PERC H740P mini monolithic card |
| 3. Card connector on the system board | 4. Card holder |

Installing the PERC H740P mini monolithic


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.

1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.
2. Open the system.
3. Align the end of the card with the card connector on the system board and the other end of the card at an angle with the plastic bracket on the system board.
4. Lower the connector side of the card into the card connector on the system board.


NOTE: Ensure that the tabs on the system board align with the screw holes on the card.


5. Using a #2 Philips screwdriver, align the screws on the card cable with the screw holes on the connector.
6. Tighten the screws to secure the card cable with the card connector on the system board.

 **NOTE:** Ensure that the screws are torqued to 5.5 lb-in (0.60N-m).

7. Close the system.
8. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.

Removing the PERC H745 adapter

 **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.

1. Turn 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 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. Disconnect any SAS cables connected to the card:
 - a. Press down and hold the metal tab on the SAS cable connector.
 - b. Pull the SAS cable out of the connector.
6. Replace the storage controller card and reconnect the SAS cable before placing them in the riser. For more information on installing the card, see [Installing the PERC H745 adapter](#).
7. Reinstall the riser on the system board and fasten the riser.
8. Close the system.
9. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.




Figure 16. Removing and installing the PERC H745 adapter

- | | |
|---------------------------------------|-----------------------------------|
| 1. Heat sink | 2. Battery |
| 3. Battery cable connector | 4. SAS/SATA backplane connector B |
| 5. SAS/SATA backplane connector A | 6. PCIe connector |
| 7. Card connector on the system board | |

Installing the PERC H745 adapter

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.

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 system board.

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 connector.
5. Connect the SAS data 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.

6. Close the system.
7. Reconnect the system to its electrical outlet and turn on the system and any attached peripherals.

Removing the PERC H745 front card

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.

1. Turn 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 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.
 - a. For systems configured with an H745 front controller in an upside down orientation it is necessary to uninstall all the drives from the backplane first and uninstall the backplane along with the controller as there is no much clearance to remove the controller without uninstalling the backplane.
5. Disconnect any SAS cables connected to the card:
 - a. Press down and hold the metal tab on the SAS cable connector.
 - b. Pull the SAS cable out of the connector.
6. Remove the PERC controller from the controller carrier.

NOTE: If the replacement storage controller is a front controller then reinsert the controller into the carrier and secure it with the appropriate screws.

7. Insert the replacement controller into the carrier and secure it with the appropriate screws.
8. Take the replacement storage controller and reconnect the SAS cable before reconnecting it to the backplane.
 - a. For systems configured with an H745 front controller in an upside down orientation and reattach the PERC controller to the backplane first before reinstalling the backplane into the system. For more information on installing the card, see [Installing the PERC H745 front controllers](#).
9. Close the system.
10. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.




Figure 17. Removing and installing the PERC H745 card

- | | |
|-----------------------------------|-----------------------------------|
| 1. PCIe cable connector | 2. Battery |
| 3. Heat sink | 4. Battery cable connector |
| 5. SAS/SATA backplane connector B | 6. SAS/SATA backplane connector A |
| 7. Power connector | |

Installing the PERC H745 front card


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.

1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.
2. Open the system.
3. Connect the PERC card to the carrier, connect the battery and ensure that the screws are properly fastened in place.
4. Align the card-edge connector with the connector on the system board.


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


5. Slide the card until it is fully seated in the connector. Tighten the carrier screws that connect to the chassis to secure the carrier.
6. Connect the SAS data 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.

7. Close the system.
8. Reconnect the system to its electrical outlet and turn on the system and any attached peripherals.


Removing the PERC H745P MX adapter card

 **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.

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

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

 **NOTE:** Perform a graceful shutdown of the sled to ensure any data in the cache is flushed 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, you must 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 SAS cable from the card. To disconnect the cable:
 - a. Press and hold the metal tab on the SAS cable connector.
 - b. Pull the SAS 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 [Installing the PERC H745P MX adapter card](#).
9. Close the sled.
10. Insert the sled into the MX chassis and turn on the system and any attached MX chassis peripherals.




Figure 18. Removing and installing the PERC H745P MX adapter

- | | |
|------------------------|---------------------------------------|
| 1. Bracket notch (3) | 2. Release lever |
| 3. Card bracket | 4. Tabs on the system |
| 5. System chassis | 6. Card connector on the system board |
| 7. PERC card connector | |

Installing the PERC H745P MX adapter card

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.

1. Turn off the sled and any attached peripherals, and remove the sled from the MX chassis.
2. Open the sled.
3. Connect the SAS 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 SAS data cable through the clip on the card and through the channel on the inner side of the chassis.
8. Attach the connector labeled "BP SAS" to connector SAS A on the backplane, and attach the connector labeled "CTRL SAS" to SAS-cable connector on the controller card.
9. Close the sled.
10. Insert the sled into the MX chassis and turn on the system and any attached MX chassis peripherals.

Removing the PERC H840 card

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.

1. Turn 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 SAS cables connected to the card:
 - a. Press down and hold the metal tab on the SAS cable connector.
 - b. Pull the SAS cable out of the connector.
5. Lift the card to remove it from the connector on the system board.
6. Replace the storage controller card and connect the cable. For more information on installing the PERC H840 card, see [Installing the PERC H840 card](#).
7. Close the system.
8. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.




Figure 19. Removing and installing the PERC H840 adapter

1. SAS cables (2)
2. SAS cable connectors (2)
3. Heat sink
4. Battery

Installing the PERC H840 card

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.

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 system board.

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 connector.
5. Connect the SAS data 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.

6. Close the system.
7. Reconnect the system to its electrical outlet and turn on the system and any attached peripherals.

Part replacement of a PERC card in eHBA mode

When replacing a PERC card, verify the controller mode of operation after replacing the card. If a controller mode change is still necessary, you may request the desired mode and reboot. Before changing the mode, ensure that all the physical drives are removed from the system. If the drives cannot be physically removed, then a force mode transition must be performed. Following reboot, foreign virtual disks should be imported.


CAUTION: The force option is for card replacement only. Use of the force option for reasons other than card replacement is not recommended.

Perform the following steps to replace the card that is operating in eHBA mode with another in RAID mode:

1. Power on the system.
2. Enter the HII configuration utility. See [Entering the HII configuration utility](#).
3. Change controller to eHBA mode. See [Manage controller mode](#).
4. Reboot the system.
5. Import any foreign virtual disks.

Part replacement of a PERC card in RAID mode with another in eHBA mode

When replacing a PERC card, verify the controller mode of operation after replacing the card. If a controller mode change is still necessary, you may request the desired mode and reboot. Before changing the mode, ensure that all the physical drives are removed from the system. If the drives cannot be physically removed, then you must perform a force mode transition. Import foreign virtual disks following reboot.

 **CAUTION: The force option is for card replacement only. Use of the force option for reasons other than card replacement is not recommended.**

Perform the following steps to replace the card that is operating in RAID mode with another in eHBA mode:

1. Power on the system.
2. Enter the HII configuration utility. See [Entering the HII configuration utility](#).
3. Change controller to RAID mode. See [Manage controller mode](#).
4. Reboot the system.
5. Import any foreign virtual disks.

Driver installation

The Dell PowerEdge RAID controller (PERC) 10 series require software drivers to operate with the supported operating systems. This chapter contains the procedures for installing the drivers for the PERC 10 cards.

NOTE: The driver for PERC 10 for VMware ESXi is packaged within the VMware ESXi ISO image downloaded from Dell. For more information, see the VMware documentation at www.dell.com/virtualizationsolutions. It is not recommended to have drivers from controllers prior to PERC 10 on the same system.

The two methods for installing a driver 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 PERC 10 family of controllers are already installed and you want to update to the latest drivers.

Topics:

- [Creating the device driver media](#)
- [Windows driver installation](#)
- [Installing the driver during a Windows Server 2012 R2 and newer installation](#)
- [Installing the driver after Windows Server 2012 R2 and newer installation](#)
- [Updating PERC 10 driver for existing Windows Server 2012 R2 and newer](#)
- [Linux driver installation](#)

Creating the device driver media

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

- [Downloading Drivers From The Dell Support Website](#)
- [Downloading Drivers From The Dell Systems Service And Diagnostic Tools Media](#)

Downloading drivers from the Dell support website

To download drivers from the Dell Support website:

1. Go to www.dell.com/support/home.
2. Enter the service tag of your system in the **Choose by Service Tag to get started** field or select **Choose from a list of all Dell products**.
3. Select the **System Type**, **Operating System**, and **Category** from the drop-down list. The drivers that are applicable to your selection are displayed.
4. Download the drivers that you require to a USB drive, CD, or DVD.
5. During the operating system installation, use the media that you created with the **Load Driver** option to load mass storage drivers. For more information on reinstalling the operating system, see the relevant section for your operating system below.

Downloading drivers from the Dell systems service and diagnostic tools media

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

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 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 this procedure for all the drivers you require.

Windows driver installation

Before you install the Windows driver for PERC 10, 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 www.dell.com/support/home.
- Create a device driver media using one of the methods listed below:
 - USB drive
 - CD
 - DVD

Installing the driver during a Windows Server 2012 R2 and newer installation

To install the driver:

1. Boot the system using the Windows Server 2012 R2, or newer media.
2. Follow the on-screen instructions until you reach **Where do you want to install Windows Server 2012 R2 or newer** window and then select **Load driver**.
3. As prompted, insert the installation media and browse to the appropriate location.
4. Select a PERC 10 series card from the list.
5. Click **Next** and continue installation.

Installing the driver after Windows Server 2012 R2 and newer installation

Perform the following steps to configure the driver for the RAID controller on a system that already has Windows installed:

1. Turn off the system.
2. Install the new RAID controller in the system.
For detailed instructions on installing the RAID controller in the system, see [Deploying the PERC card](#).
3. Turn 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 from the **Locate Driver Files** screen.
7. Click **Next**.
The wizard detects and installs the appropriate device drivers for the new RAID controller.
8. Click **Finish** to complete the installation.
9. Reboot the system when prompted.

Updating PERC 10 driver for existing Windows Server 2012 R2 and newer

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

1. Insert the CD, DVD, or USB drive media containing the driver.
2. Select **Start > Settings > Control Panel > System**.
The **System Properties** screen is displayed.
NOTE: The path to **System** might vary depending 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** might vary depending on the operating system family.
5. Expand **SCSI and RAID Controllers** by double-clicking the entry or by clicking on the plus symbol next to **SCSI and RAID Controller**.
NOTE: In Windows Server 2012 R2 and newer, the PERC 10 series of cards are listed under **Storage Controllers**.
6. Double-click the RAID 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 CD, DVD, or USB 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 take place.
NOTE: Dell provides the Dell Update Package (DUP) to update drivers on systems running Windows Server 2012 R2 and newer operating system. DUP is an executable application that updates drivers for specific devices. DUP supports command line interface and silent execution. For more information, see <https://www.dell.com/support>.

Linux driver installation

- NOTE:** PERC 10 series of drivers support PERC 9 family controllers and do not require separate driver installations.
- NOTE:** The driver update disk (DUD) images are created only for those operating system releases in which the native (in-box) driver is insufficient for installation. In the event that an operating system is being installed with a corresponding DUD image, follow the instructions below.
- NOTE:** To view the complete list of boot loader options, see the installation guide of your operating system.
- NOTE:** If using out-of-box drivers with RHEL 7 and higher, a tainted kernel message will be displayed in the log. RedHat does not provide a mechanism to sign external drivers for RHEL. This leads to kernel tainting whenever external drivers are loaded in all the versions. Kernel messages can be

```
megaraid_sas>Loading out-of-tree module taints kernel
```

```
megaraid_sas:module verification failed: signature and/or required key missing-tainting kernel
```

Installing or updating the RPM driver package with KMOD support

NOTE: This procedure is applicable for Red Hat Enterprise Linux 7.x.

Perform the following steps to install the RPM package with KMOD support:

1. Uncompress the gzipped tarball driver release package.
2. Install the driver package using the command: `rpm -ihv kmodmegaraid_ sas-<version>.rpm.`

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

3. If the previous device driver is in use, you must reboot the system for the updated driver to take effect.
4. Verify that the driver is loaded with the following system commands: `modinfo megaraid_ sas.`

Installing or updating the RPM driver package with KMP support

NOTE: This procedure is applicable for SUSE Enterprise Linux 12.x.

Perform the following steps to install the RPM package with KMP support:

1. Uncompress the gzipped tarball driver release package.
2. Install the driver package using the command: `rpm -ihv kmpmegaraid_ sas- <version>.rpm.`

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

3. If the previous device driver is in use, you must reboot the system for the updated driver to take effect.
4. Verify that the driver is loaded with the following system commands: `modinfo megaraid_ sas.`

Loading the driver during operating system installation

1. Perform the following operation to install the driver media:
 - a. Download the PERC linux driver ISO, or install the LC driver pack.
 - b. Mount the ISO to the Server, burn the ISO to CD/DVD, or copy the ISO file to USB. The USB has to match with the ISO.
 - c. For LC driver pack, boot the life-cycle controller and go through the operating system deployment wizard.
2. Boot to the installer.
3. In the Installation screen, press E.
4. Perform the following operation:

- If the operating system is Red Hat Enterprise Linux 7 or RHEL 8, 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 12/15, 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.
A screen is displayed prompting you to select the driver media (USB, CD, ISO, and so on).
7. When prompted select the driver media.
If applicable select the PERC driver `...megaraid_ sas...`

NOTE: Ensure that the driver is selected with an X symbol.

8. The driver should be extracted or loaded.
9. Before proceeding or exiting 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.

10. Press C or exit to go to the installation.

Firmware

This section provides information about downloading and installing the firmware using Dell Update Package (DUP).

Topics:

- [Installing the firmware using DUP](#)

Installing the firmware using DUP

1. Navigate to www.dell.com/support/home.
2. Locate your controller.
3. Download the DUP.
 - a. For Window/iDRAC update, download Windows executable file.
 - b. For Linux update, download `.bin` file.
4. Install the DUP.
 - a. For Windows, run the executable in Windows environment.
 - b. For Linux, run `.bin` file in Linux environment.
 - c. For iDRAC, navigate to **system iDRAC > Maintenance > System Update**, upload Windows executable, and then install.

NOTE: For VMware, firmware should be updated through iDRAC or the PERC CLI utility.

NOTE:

- Firmware downgrade from version 50.3.0-1022 or later to the older version is not supported via DUP.
- Firmware downgrade from version 50.5.0-1750 to earlier versions is not supported while the controller is in eHBA mode.
- Firmware downgrade from version 50.5.1-2633 to earlier versions is not supported while the controller is in OpenManage Secure Enterprise Key Manager mode.
- PERC H345 and H745 supports firmware version 50.9.1-2905 or later.
- PERC 745P MX supports firmware version 50.3.0-1512 or later.

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 RAID disk groups, virtual disks, and physical disks. This utility is independent of the operating system.

Topics:

- [Entering the HII configuration utility](#)
- [Exiting the HII configuration utility](#)
- [Navigating to Dell PERC 10 configuration utility](#)
- [HII Configuration utility dashboard view options](#)
- [Configuration management](#)
- [Controller management](#)
- [Virtual disk management](#)
- [Physical disk management](#)
- [Hardware components](#)
- [Security key management in the HII configuration utility](#)

Entering the HII configuration utility

Perform the following steps to boot to the HII configuration utility:

1. Turn on the system.
2. While the system startup, press <F2> to enter **System Setup**.
3. Click **Device Settings**.
Device Settings screen lists all the RAID controllers in the system.

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

NOTE: For more information in all the options, click Help that is available on the top right-hand corner of the browser screen. Help information for individual option menus can also be viewed by scrolling down on each option.

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.

Exiting the HII configuration utility

To exit the HII configuration utility, perform the following steps:

1. Click **Finish** at the bottom-right corner on the **System Setup Main Menu** screen.
Displays a warning message to confirm your choice.
2. Click **Yes** to exit the HII configuration utility.

Navigating to Dell PERC 10 configuration utility

1. Enter the UEFI configuration Utility. See [Entering the HII configuration utility](#).
The **Device Settings** screen displays a list of NIC ports and the Dell PERC 10 configuration utility.
2. To enter PERC 10 configuration utility, click **Dell PERC 10 Configuration Utility**.
The **Dashboard view** screen is displayed.

HII Configuration utility dashboard view options

The first screen that is displayed when you access the HII Configuration Utility is the **Dashboard View** screen. The following table provides detailed information about the options available on the **Dashboard View** screen.


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• Physical Disk Management• Hardware Components
Help	Provides context sensitive help message.
Properties	Displays the following information about the controller: <ul style="list-style-type: none">• Status — displays the status of the controller.• Backplane — displays information about the number of backplanes connected to the controller.• BBU — displays information about the availability of Battery Backup Unit (BBU).• Enclosure — displays information about the number of enclosures connected to the controller.• Physical Disks — displays information about the number of physical disks connected to the controller.• Disk Groups — displays information about the number of disk groups connected to the controller.• Virtual Disks — displays information about the number of virtual disks connected to the controller.
View server profile	Displays HII Spec version supported on the system and also displays the following menu options for controller components: <ul style="list-style-type: none">• Controller Management• Hardware Components• Physical Disk Management• Virtual Disk Management
Actions	Displays the following options: <ul style="list-style-type: none">• Configure — displays configuration options that are supported by the controller.• Set Factory Defaults — restore factory default values for all controller properties.
Background operations	Displays if virtual disk or physical disk operations are in progress.

Configuration management

Auto Configure RAID 0

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Configuration Management > Auto Configure RAID 0**.
3. Select **Confirm** and click **Yes** to continue.
A RAID 0 Virtual disk is created on all physical disks that are in Ready state.

 **NOTE:** This feature is not supported on PERC H345 and eHBA mode.

Creating virtual disks

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Configuration Management > Create Virtual Disk**.
The following list of options are displayed for you to define the virtual disk parameters:
 - Select RAID level — allows you to choose the RAID level of your choice.
 - Secure Virtual Disk — If you want to create a secured virtual disk, select **Secure Virtual Disk**.

NOTE: The Secure Virtual Disk option is enabled by default, only if the security key has been configured. Only SED physical disks are listed.
 - Select Physical Disks From — allows you to select one of the physical disk capacities:
 - **Unconfigured Capacity:** creates a virtual disk on unconfigured physical disks.
 - **Free Capacity:** utilizes unused physical disk capacity that is already part of a disk group.
 - Select Physical Disks — If you want to select the physical disks from which the virtual disks are being created, click **Select Physical Disks**. This option is displayed if you select **Unconfigured Capacity** as your physical disk capacity.
 - Select Disk Groups — If you want to select the disk groups from which the virtual disks are being created, click **Select Disk Group**. This option is displayed if you select **Free Capacity** as your physical disk capacity.
 - Configure Virtual Disk Parameters — allows you to set the virtual disk parameters when creating the virtual disk. For more information, see [Configuring virtual disk parameters](#).
3. Click **Create Virtual Disk**.
The virtual disk is created successfully.

Configuring virtual disk parameters

1. Create a virtual disk, see [Creating the virtual disks](#).
The **Configure Virtual Disk Parameters** section is displayed on the **Create Virtual Disk** screen.
2. In the **Configure Virtual Disk Parameters** section, you can set the following virtual disk parameters:

Table 8. Configuring virtual disk parameters

Virtual disk parameters	Description
Virtual Disk Name	Allows you to enter the name for the virtual disk <div style="border-left: 1px solid black; padding-left: 5px; margin-left: 20px;"> <p>NOTE: Allowed characters are A-Z, a-z, 0-9, underscore (_), and hyphen (-) only.</p> </div>
Virtual Disk Size	Displays the maximum capacity available for the virtual disk
Virtual Disk Size Unit	Displays the virtual disk storage space in megabytes, gigabytes, and terabyte
Strip Element Size	Allows you to select the strip element size The 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 By default, the strip element size is set to 256 KB.
Read Policy	Displays the controller read policy You can set the read policy to: <ul style="list-style-type: none"> • No read ahead—specifies that the controller does not use read ahead for the current virtual disk. • Read ahead—specifies that the controller uses read ahead for the current virtual disk. Read ahead capability allows the controller to read sequentially ahead of requested data and store the additional data in the cache memory, anticipating that the data is required soon. By default, the read cache policy is set to read ahead.
Write 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. By default, the write policy is set to Write Back.
Disk Cache	Allows you to set the disk cache policy to default, enable, or disable By default, the disk cache is set to default.

Table 8. Configuring virtual disk parameters (continued)

Virtual disk parameters	Description
Default Initialization	<p>Displays the virtual disk initialization options. You can set the default initialization to:</p> <ul style="list-style-type: none">• No — The virtual disk is not initialized.• Fast — The first 8 MB of the virtual disk is initialized.• Full — The entire virtual disk is initialized. <p>For more information, see Virtual disk initialization. By default, the default initialization is set to No.</p>

Creating profile based virtual disk

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 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
3. Based on the RAID mode selected, one or more the physical disk selection criteria is displayed.
4. From the **Physical Disk Selection Criteria** drop-down box, select a criterion based your requirement.
The Profile Parameters of the selected option is displayed.
5. Click **Create Virtual Disk**.
6. Select **Confirm** and click **Yes** to continue.
The virtual disk is created with the parameters of the profile selected.

Viewing disk group properties

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Configuration Management > View Disk Group Properties**.
The list of disk group properties are displayed:
 - Capacity Allocation — displays all the virtual disks associated with the specific disk group. It also provides information about the available free space.
 - Secured — displays whether the disk group is secured or not.

Deleting configurations

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Configuration Management > Clear Configuration**.
A screen is displayed asking if you are sure you want to perform the operation.
3. Select **Confirm** and click **Yes** to continue.
The virtual disks and hot spare disks available on the controller are deleted successfully.

Controller management

Clearing controller events

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).

2. Click **Main Menu > Controller Management > Advanced Controller Management**.
3. Click **Clear Controller Events**.
A screen is displayed indicating that the operation is completed successfully.
4. Click **Ok**.

Saving controller events


1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Management**.
3. Click **Save Controller Events**.
A screen is displayed indicating that the operation is completed successfully.
4. Click **Ok**.

Saving debug log

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Management**.
3. Click **Save Debug Log**.
A screen is displayed indicating that the operation is completed successfully.
4. Click **Ok**.

Managing link speed


1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Management**.
3. Click **Manage Link Speed**.
You can view/change the physical disk settings for the controller, if required. The possible options are:
 - Auto
 - 3 Gbps
 - 6 Gbps
 - 12 Gbps
4. Click **Ok** to save the changes.

 **NOTE:** To reflect the changes, perform an AC power cycle that is disconnecting and reconnecting the power source.

Restoring factory default settings

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Set Factory Defaults**.
A screen is displayed asking you to confirm the operation.
3. Select **Confirm** and click **Yes** to continue.

Manage controller mode


 **NOTE:** All existing configurations on the controller must be deleted before proceeding with a controller mode change. Once the mode transition operation is confirmed, the process can take up to 5 seconds to complete; do not power off or reset the system until the mode transition operation is complete. This action is not supported on PERC H345 controller.

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Management > Manage Controller Mode**.
3. Click **Switch to Enhanced HBA mode** or **Switch to RAID mode**.

 **CAUTION:** Only use force switch controller mode as part of a controller replacement operation. Consult with technical support for assistance.

4. Check the box to **Confirm** the controller mode change.
5. Click **Yes** to confirm the selection.
6. Click **OK** to acknowledge the change.
A reboot is required to complete the mode transition. Until this reboot is performed, avoid requesting any further changes. Before the next reboot:
 - All configuration commands are blocked
 - Controller firmware flash is not allowed


Following the reboot, the mode transition is complete and the controller resumes normal operation.

 **NOTE:** Mode change from eHBA to RAID may require a cold boot for Management Applications to properly inventory drives.


Changing profile using HII Configuration Utility


Perform the following steps to change the profile to HII configuration utility:

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Management > Manage Controller Profiles**. Current profile and profile properties are displayed.
3. Change profile using the **Choose Profile** option.

 **NOTE:** By default, eHBA mode only supports profile 23_PD240.

4. Select **Set Profile**.
Click **Reboot**.

 **NOTE:** For the changes to reflect reboot is required.

 **NOTE:** Profile change fails:

- if new profile supports less number of drives than the number of drives in the current topology.
- if background operations (rebuild, copy back, full initialization, Background Initialization, Patrol Read, CC) are active.
- if background operations start after profile change but before the system is rebooted.

Advanced controller properties

Clearing the cache memory

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 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.

Setting patrol read

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 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.
4. Set the **Mode** to **Auto**, **Manual**, or **Disabled**.

5. Click **Apply Changes**.

Enabling physical disk power management

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Properties**.
3. Click **Physical Disk Power Management**.
The following list of options is displayed:
 - Time Interval for Spin Down — allows the user to specify the delay time before a disk is spun down.
 - Spin Down Hot Spare — allows you to enable or disable the spin down of hot spare disks.
 - Spin Down Unconfigured Good — spin down of un-configured disks.
4. Select the applicable options and click **Apply Changes**.
The changes made are saved successfully.

Configuring hot spare

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Properties**.
3. Click **Spare**.
The following list of options are displayed:
 - Persistent Hot Spare — allows you to enable or disable the ability to have same system backplane or storage enclosure disk slots dedicated as hot spare slots.
 - Allow Replace Member with Revertible Hot Spare — allows you to enable or disable the option to copy the data form a hot spare disk to physical disk.
 - Auto Replace Member on Predictive Failure — allows you to enable or disable the option to start a Replace Member operation if a predictive failure error is detected on a physical disk.
4. Select the applicable option and click **Apply Changes**.
The changes made are saved successfully.

Setting task rates

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Properties**.
3. Click **Task Rates**.
The following options are displayed:
 - Background Initialization (BGI) Rate
 - Consistency Check Rate
 - Rebuild Rate
 - Reconstruction Rate
4. You can make the necessary changes and then click **Apply Changes**.
The task rates operation is completely successfully.

Enabling auto import

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Properties**.
3. In the **Controller Properties** section, set the **Auto Import Foreign Configuration** option to **Enabled**.
4. Click **Apply Changes**.
The auto import is enabled successfully.

Disabling auto import

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Properties**.

3. In the **Controller Properties** section, set the **Auto Import Foreign Configuration** option to **Disabled**.
4. Click **Apply Changes**.
The auto import is disabled successfully.

Selecting boot mode

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Properties**.
3. In the **CONTROLLER PROPERTIES** section, select boot mode from the **Boot Mode** drop-down box.
The following lists of boot mode options appear:

Table 9. Boot mode options

Option	Description
Stop on errors	The system stops during boot for errors which require attention from the user to rectify the issue.
Pause on errors	System pauses during boot to show errors but continue boot after it times out. Only critical events with an infinite timeout halt boot and require the user's attention to correct the issue.
Ignore errors (headless mode)	System ignores most errors during boot in headless mode and prevents the system from halting during boot. Headless mode is designed for systems that do not have a dedicated monitor or a peripheral connected, requiring the user to manage errors during boot. In headless mode, errors can be monitored by viewing PERC events, the life-cycle log, operating system logs, and the system event log. NOTE: In headless mode, any foreign disk is automatically imported during boot.
Safe mode on errors	System is routed to safe mode when critical errors arise. PERC firmware disables most of the features on the controller, and the controller requires attention from the user to rectify the issue.

NOTE: In UEFI BIOS mode, errors with timeouts do not appear during boot. It is designed to arise only in legacy BIOS mode.

NOTE: By default, the boot mode option is set to pause on errors.

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

Aborting the consistency check

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 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**.
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.

Enabling boot support

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Controller Management**.
3. From the **Select Boot Device** drop-down box, select the primary bootable device.

NOTE: **Select Boot Device** is only applicable in legacy BIOS mode.

NOTE: 4K boot support is only available in UEFI mode and managed by the boot loader.

NOTE: In **Select Boot Device**, you will not be able to view 4K devices. To view all the virtual disks created, navigate to the **Virtual Disk Management** screen in HII. For more information, see [Virtual disk management](#).

NOTE: If no boot device is selected, the first virtual disk will be set as the boot device on the next reboot. A Non-RAID disk will not be auto-selected as the boot device.

- Click **Apply Changes**.
Boot support is enabled for the selected controller.

Virtual disk management

Virtual disk numbering

Virtual disks are numbered in ascending order beginning with zero in RAID mode, whereas in eHBA mode they are numbered in descending order beginning with 239.

Viewing virtual disk properties

- Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
- Click **Main Menu > Virtual Disk Management**.
The virtual disks associated with the RAID controller are displayed along with their target ID, name, RAID level, size, and state.
- To view the properties, click on the virtual disk. You can view the following properties of the virtual disk:

Table 10. Virtual disk properties

Option	Description
Operation	List of operations you can perform on the selected virtual disk. The options are: <ul style="list-style-type: none"> Blink Unblink Delete Virtual Disk Reconfigure Virtual Disks Fast Initialization Slow Initialization Secure Virtual Disk NOTE: The secure virtual disk option is only supported on SED physical disks. The security key must be configured to enable the feature. Expand Virtual Disk Consistency Check
Name	Indicates the name of the virtual disk.
Status property	Indicates the status of the virtual disk. The possible options are: <ul style="list-style-type: none"> Optimal Degraded Offline Failed
RAID level	Indicates the RAID level of the virtual disk.
Size	Indicates the size of the virtual disk.

- Click **Advanced...**
You can view the following additional properties of the virtual disk:

Table 11. Advanced properties of the virtual disk

Option	Description
Logical sector size	Indicates the logical sector size of this virtual disk.
Strip element size	Indicates the strip element size for the virtual disk.
Status	Indicates the status of the virtual disk.
Secured	Indicates whether the virtual disk is secured or not.
Bad blocks	Indicates whether the virtual disk has corrupted blocks.

Viewing physical disks associated with a virtual disk

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 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.

Configuring virtual disk policies

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Virtual Disk Management**.
All the virtual disks associated with the RAID controller are displayed.
3. Click **Advanced...**
You can view the following virtual disk policies:

Table 12. Virtual disk policies

Option	Description
Current write cache	Indicates the current write cache policy for the virtual disk.
Default write cache	Allows selection of the write cache policy for the virtual disk. The possible options are: <ul style="list-style-type: none">• Write Through• Write Back• Force Write Back
Read cache policy	Allows selection of the read cache policy for the virtual disk. The possible options are: <ul style="list-style-type: none">• No Read Ahead• Read Ahead
Disk cache	Allows selection of the disk cache policy for the virtual disk. The possible options are: <ul style="list-style-type: none">• Default (Disk Default)• Enable• Disable

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

Configuring Virtual Disks


When configuring the virtual disks, you should consider the workload intended; RAID1: for simple boot disk; RAID5 or 6: for file or web servers (sequential reads/writes of files); RAID10: for transactional database (small random reads and writes).


Virtual disks configured on hard drives should use the controller default cache setting of Write Back and Read Ahead.

Virtual disks configured on SSDs can use the same controller defaults settings as hard drives. Most users perform a copy of OS files or a data base to the new array. This setting provides optimum performance in this configuration.

Once the copy is complete, the array can be used as it is depending on the number and type of SSDs. It is recommended to enable FastPath by changing the controller's Write cache policy to Write Through and the Read cache policy to No Read Ahead. FastPath is developed to achieve the best random read/write performance from SSDs.

Only IO block sizes smaller than the virtual disk's stripe size are eligible for FastPath. In addition, there should be no background operations (rebuild, initialization) running on the virtual disks. FastPath is disabled if there is active background operation.


 **NOTE:** RAID 50, and RAID 60 virtual disks cannot use FastPath.

 **NOTE:** The Physical Disk Power Management feature is not applicable to FastPath-capable virtual disks.

Performing expand virtual disk operation

To enable expand virtual disk feature from the HII Configuration Utility, perform the following steps:

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Virtual Disk Management**.
The list of virtual disks is displayed.
3. Select the virtual disk.
4. From the **Operations** drop-down menu, select **Expand Virtual Disk**.


 **NOTE:** You can view the Expand Virtual Disk feature only if there is free space available in the associated disk group.

5. Click **Go**.
6. To expand virtual disk, enter the percentage of available capacity, and then click **Ok**.
A screen is displayed asking if you are sure you want to perform the operation.
7. Select the **Confirm** option.
8. Click **Yes**.
The expand virtual disk operation is completed successfully.

Performing consistency check

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

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 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** drop-down menu, select **Check Consistency**.
5. Click **Go**.
A screen is displayed asking if you are sure you want to perform the operation.
6. Select the **Confirm** option.
7. Click **Yes**.
The consistency check operation is completed successfully.

Physical disk management

Viewing physical disk properties

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Physical Disk Management**.
All the physical disks associated with the RAID controller are displayed.
3. To view the properties, click on the physical disk. The following properties can be viewed on the physical disk:

Table 13. Physical disk properties


Option	Description
Operation	List of operations you can perform on the selected physical disk. The options are: <ul style="list-style-type: none"> • Blink • Unblink • Force Offline • Replace Member • Erase Options • Global and Dedicated hotspare • Convert to RAID Capable • Convert to Non-RAID Disk
Device ID	Unique identifier of the physical disk.
Connection	Number of paths and how wide the paths are to the physical disk. Indicated if a physical disk is in a multipath configuration for H745P MX and H840.
Backplane ID	Backplane ID in which the physical disk is located in for H345, H740P, H740 mini, H745, H745P MX, and H840.
Enclosure position	The enclosure in which the physical disk is located for H840.
Chassis slot value	The chassis slot value of the chassis slot in which the physical disk is located when the physical disk is in a MX 5016s storage sled for the controller H745P MX.
Slot number	The drive bay in which the physical disk is located for the corresponding backplane or enclosure to which the controller is connected.
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 disk.
Serial number	Serial of the physical disk.
Manufacturing date	Date on which the physical disk has been manufactured.
Associated virtual disks	Virtual disks associated with the physical disk.


4. Click **Advanced...**
You can view the following additional properties of the physical disk:

Table 14. Advanced physical disk properties

Option	Description
Logical sector size	Logical sector size of the selected physical disk.
Physical sector size	Physical sector size of the selected physical disk.

Table 14. Advanced physical disk properties (continued)

Option	Description
SMART status	SMART status of a physical disk.
Revision	Firmware version of the physical disk.
SAS address	SAS address of the physical disk.
Physical disk power state	Power condition (On or Power Save) of the physical disk.
Disk cache settings	Disk cache settings.  NOTE: Disk cache for SATA Gen3 drives is disabled by default.
Hard disk drive RPM	RPM of the hard drive.
Available space	Available size of the physical disk.
Used space	Configured space of the physical disk.
Disk protocol	Type of hard disk used.
Device speed	Speed of the physical disk.
Negotiated link speed	Negotiated link speed of the device.
Encryption capable	Encryption capability of the physical disk.
Secured	Security status of the physical disk.
Cryptographic erase capable	Cryptographic erase capability of the physical disk.

 **NOTE:** When you select the **Patrol Read** option in the **Advanced...** section, an error message may be displayed. The error message can be safely ignored.


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.

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

The Cryptographic erase feature is supported only on Instant Secure Erase (ISE) and Self Encrypting Drives (SED) drives.

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Physical Disk Management**.
The list of physical disks is displayed.
3. Select a physical disk.
4. From the **Operations** drop-down menu, select **Cryptographic Erase**.

 **NOTE:** If the drive installed is ISE or SED capable, only then the Cryptographic erase option is displayed.

5. Click **Go**.
A screen is displayed asking if you are sure you want to perform the operation.
6. Select the **Confirm** option.
7. Click **Yes**.
The Cryptographic erase operation is completed successfully.

Physical disk erase

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

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Physical Disk Management**.

The list of physical disks is displayed.

3. Select a physical disk.
4. From the **Operations** drop-down menu, select **Physical Disk Erase**.

 **NOTE:** If the drive installed is neither SED or ISE capable, then only the Physical Disk Erase option is displayed.

5. Click **Go**.
A screen is displayed asking if you are sure you want to perform the operation.
6. Select the **Confirm** option.
7. Click **Yes**.
The physical disk erase operation is completed successfully.

Assigning global hot spare

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

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Physical Disk Management**.
The list of physical disks is displayed.
3. Select the physical disk.
4. From the **Operations** drop-down menu, select **Assign Global Hot Spare**.
5. Click **Go**.
A screen is displayed asking if you are sure you want to perform the operation.
6. Select the **Confirm** option.
7. Click **Yes**.
The global hot spare disk is created successfully.

Assigning dedicated hot spare

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

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Physical Disk Management**.
The list of physical disks is displayed.
3. Select the physical disk.
4. From the **Operations** drop-down menu, select **Assign Dedicated Hot Spare**.
5. Click **Go**.
A screen is displayed asking if you are sure you want to perform the operation.
6. Select the **Confirm** option.
7. Click **Yes**.
The dedicated hot spare disk is created successfully.

Convert to RAID capable

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

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Physical Disk Management**.
The list of physical disks appears.
3. Select the physical disk.
4. From the **Operations** drop-down menu, 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 dedicated hot spare disk is created successfully.

Convert to Non-RAID disk

To assign a disk from the HII Configuration Utility, perform the following steps:

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Physical Disk Management**.
The list of physical disks appears.
3. Select the physical disk.
4. From the **Operations** drop-down menu, select **Convert to Non-Raid disk**.
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 disk is created successfully.

Hardware components

Viewing battery properties

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Hardware Components > Battery Management**.
The battery and capacity information are displayed.
3. You can view the following properties of the battery:

Table 15. Battery properties

Field	Description
Type	Displays the type of battery available.
Status	Displays the current status of the battery.
Temperature	Displays the current temperature of the battery and also indicates whether the temperature is normal or high.
Charge	Displays the available charge of the battery in percentage.

4. Displays click **Advanced...**.
The additional advanced properties of the physical battery are displayed.
5. You can view the following advanced properties of the battery:

Table 16. Advanced battery properties

Field	Description
Status	Displays whether the current status of the battery is learning, degraded, or failed.
Voltage	Displays whether the voltage status of the battery is normal or high.
Current	Displays power consumption of the battery in milliamps (mA).
Full capacity	Displays the maximum charge capacity of the battery.
Remaining capacity	Displays the current charge capacity of the battery.
Expected margin of error	Displays expected margin of error.
Completed discharge cycles	Displays the completed discharge cycles.
Learn mode	Displays the condition of the battery. The learn cycle is a periodic operation that calculates the charge that is remaining in the battery to ensure there is sufficient energy.

Viewing physical disks associated with an enclosure

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Hardware Components > Enclosure Management**.
3. From the **Select Enclosure** field, choose the enclosure for which you need to view the physical disks. All the physical disks that are associated with the virtual disk are displayed.
4. Click the **Attached Physical Disks** drop-down box. All the physical disks that are associated with the selected enclosure are displayed.

Security key management in the 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 provide detailed instructions to perform the configuration tasks. The contents in the following section apply to the **HII Configuration Utility**. For more information on the management applications, see [Management applications for PERC cards](#).

- 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**—Creates, changes, or deletes the security settings on a controller.
- 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 **Physical Disk Management** screen displays physical disk information and action menus. You can perform the following security-related actions through the physical disk management menu:
 - **Cryptographic Erase**—Permanently erases all data on the physical disk and resets the security attributes.

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

Security key and RAID management

Topics:

- Security key implementation
- Local Key Management
- Creating security key
- Changing Security Settings
- Disabling security key
- Create secured virtual disk
- Secure non-RAID disks
- Secure preexisting virtual disk
- Importing secure virtual disk
- Import secured non-RAID disk
- Dell EMC OpenManage Secure Enterprise Key Manager


Security key implementation

The Dell PowerEdge RAID Controller (PERC) 10 series of cards support Self-Encrypting Disks (SED) 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 under Local Key Management (LKM). The key can be escrowed in to 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. In order to take advantage of this feature, you must:


1. Have SEDs in your system.
2. Create a security key.

Local Key Management


You can use Local Key Management (LKM) to generate the key ID and the passphrase required to secure the virtual disk. You can secure virtual disks, change security keys, and manage foreign configurations using this security mode.

 **NOTE:** Under LKM, you are prompted for a passphrase when you create the key. This mode is not supported on PERC H345.


Creating security key

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


1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Controller Management > Advanced Controller Management > Enable Security**.
3. Select the **Security Key Management** mode as **Local Key Management**.
4. Click **Ok**.
5. In the **Security Key Identifier** field, enter an identifier for your security key.

 **NOTE:** The Security Key Identifier is a user supplied clear text label used to associate the correct security key with the controller.

6. If you want to use the passphrase generated by the controller, click **Suggest Passphrase**. Assigns a passphrase suggested by the controller automatically.
7. In the **Passphrase** field, enter the passphrase.

 **NOTE:** Passphrase is case-sensitive. You must enter minimum 8 or maximum 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** field, re-enter the passphrase to confirm.

 **NOTE:** If the Passphrase entered in the Passphrase and Confirm fields do not match, then you are prompted with an error message to enter the passphrase again.

9. Select the **I recorded the Security Settings for Future Reference** option.

10. Click **Enable Security**.

The Security Key is created successfully.

Changing Security Settings

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).

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

3. Select security identifier:

a. To change the **Security key Identifier** enter a new key identifier in **Enter a New Security Key identifier** text box.

b. To keep existing key identifier, select **Use the existing Security Key Identifier** check box.

4. Enter the existing passphrase.

5. Set passphrase:

a. To change the security passphrase, enter a new passphrase in the **Enter a New Passphrase** text box. Re-enter the new passphrase to confirm.

b. To keep the existing passphrase, select **Use the existing passphrase**.


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

7. Click **Save Security Settings**.

8. Select **Confirm** and then click **Yes**.

Security settings changed successfully.

Disabling security key

 **NOTE:** Disabling Security Key is active if there is a security key present on the controller.

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).


2. Click **Main Menu > Controller Management > Advanced Controller Management > Disable Security**.

You are prompted to confirm whether you want to continue.

3. Select the **Confirm** option.

4. Click **Yes**.


The security key is disabled successfully.

 **NOTE:** All virtual disks must be deleted or removed to disable security.

 **WARNING:** Any un-configured secured disks in the system will be repurposed.

Create secured virtual disk

To create a secured virtual disk, the controller must have a security key established first. See [Creating security key](#).

 **NOTE:** Combining SAS and SATA hard drives within a virtual disk is not supported. Also, combining hard drives and Solid-State Drives (SSDs) within a virtual disk is not supported.

After the security key is established, perform the following steps:

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).

2. Click **Main Menu > Configuration Management > Create Virtual Disk**.

For more information, see [Creating the virtual disks](#).

3. Select the **Secure Virtual Disk** option.
4. Click **Create Virtual Disk**.
The secure virtual disk is created successfully.


Secure non-RAID disks

In HII, secure non-RAID disks by using the security settings of the controller.

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Physical Disk Management**.
The list of virtual disks is displayed.
3. Select a non-RAID disk.
4. From the **Operations** drop-down menu, select **Secure Non-RAID Disk**.

Secure preexisting virtual disk

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 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.


Importing secure virtual disk

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


1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Configuration Management > Manage Foreign Configurations > Preview Foreign Configurations**.
3. Click **Import Foreign Configuration**.
A screen is displayed asking if you are sure you want to perform the operation.
4. Enter **Passphrase** if importing virtual disk with a different passphrase.
5. Select the **Confirm** option.
6. Click **Yes**.
The foreign configuration is imported successfully.

Import secured non-RAID disk

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




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

1. Enter the **Dell PERC 10 Configuration Utility**. See [Navigating to Dell PERC 10 configuration utility](#).
2. Click **Main Menu > Configuration Management > Manage Foreign Configurations**.
3. Click **Enter Passphrase for Locked Disks**.
A screen is displayed asking if you are sure you want to perform the operation.
4. Enter **Passphrase** if importing disk with a different passphrase.
5. Select the **Confirm** option.
6. Click **Yes**.

 **NOTE:** If **Auto-Configure** for non-RAID disk is enabled, the disk becomes a non-RAID disk. Else, it is unconfigured.

Dell EMC OpenManage Secure Enterprise Key Manager

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 entire system is stolen. Refer to the www.dell.com/idracmanuals for more information on configuring OpenManage Secure Enterprise Key Manager, as well as Secure Sockets Layer (SSL) or Transport Layer Security (TLS) related configuration.

-  **NOTE:** Downgrade of PERC firmware to a firmware that does not support enterprise key management while enterprise key manager mode is enabled, is blocked.
-  **NOTE:** When replacing a controller enabled with enterprise key management, lifecycle controller part replacement will re-configure the new controller to match the existing controller's configuration.
-  **NOTE:** If key exchange fails during boot, view and correct any connection issues with the key server identified in the iDRAC lifecycle log. Then the system can be cold booted.





Supported controllers for OpenManage Secure Enterprise Key Manager

Enterprise key manager mode is supported on the PERC H740P adapter, PERC H740P mini adapter, H745 front card, H745 adapter and on split backplane mode. For more information on supported platforms, see www.dell.com/idracmanuals.

Enterprise key manager mode is not supported on the PERC H345, PERC H840 external adapter, PERC H745P (NGM), or on the H740P and H745 with eHBA mode enabled.

Managing enterprise key manager mode

Enterprise key manager features are managed by iDRAC. For instructions on enabling enterprise key manager mode, see www.dell.com/idracmanuals.

-  **NOTE:** If preserved cache is present, the controller will not allow OpenManage Secure Enterprise Key Manager mode to be enabled.
-  **NOTE:** Transitioning a controller from LKM mode to enterprise key manager mode without disabling security or the reverse is not supported.
-  **NOTE:** When enterprise key manager mode is enabled, the controller waits up to two minutes for iDRAC to send keys, after which the PERC continues to boot.
-  **NOTE:** Rotation of keys is done by iDRAC. Any attempt to re-key the controller through a different management application is not supported.

Disabling enterprise key manager mode

Enterprise key manager mode can be disabled from any supported [Management applications for PERC cards](#). Enterprise key manager mode can be disabled from any supported [Management applications for PERC cards](#).

Managing virtual disks in enterprise key manager mode

Virtual disks are managed in the same way in enterprise key manager mode as in local key manager mode. SED capable virtual disks can be secured during or after creation. See [Creating secured virtual disk](#).

Migration of drives from local key management to enterprise key management

Local key management drives can be migrated to an enterprise key management enabled system, but the controller cannot be transitioned from local key management mode to enterprise key manager mode or the reverse without first disabling security on the controller. Perform the following steps to migrate local key management drives to enterprise key management:

1. Save the current local key management security key.
2. Shut down both systems.
3. Remove the local key management drives and re-insert them to the enterprise key manager enabled system.
4. Power on the enterprise key manager system.
5. Navigate to HII foreign configuration.
6. Enter the local key management keys for those drives.
7. Import the configuration.

NOTE: Once local key management drives are migrated to enterprise key manager, they cannot be migrated back to local key management mode. The drives have to be cryptographically erased to disable security and then converted back to local key management disks. For more information on performing this action, contact <https://www.dell.com/supportassist>.

Troubleshooting

To get help with your Dell PowerEdge RAID Controller (PERC) 10 series, you can contact your Dell Technical Service representative or see <https://www.dell.com/support>.

Topics:

- [Single virtual disk performance or latency in hypervisor configurations](#)
- [Adapter at baseport not responding error message](#)
- [BIOS disabled error message](#)
- [Configured disks removed or not accessible error message](#)
- [Dirty cache data error message](#)
- [Discovery error message](#)
- [Drive Configuration Changes Error Message](#)
- [Windows operating system installation errors](#)
- [Firmware fault state error message](#)
- [Extra enclosure error message](#)
- [Foreign configuration found error message](#)
- [Foreign configuration not found in HII error message](#)
- [Degraded state of virtual disks](#)
- [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](#)
- [HII error messages](#)
- [System reports more drive slots than what is available](#)

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, read ahead 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 reconstructions are running on the virtual disk. See your hypervisor storage best practices or performance best practices guides for additional configuration support.

Adapter at baseport not responding error message

Error Message: Adapter at Baseport xxxx is not responding, where xxxx is the baseport of the controller.

Corrective Action: Contact Dell Technical Support.

BIOS disabled error message

Error Message: BIOS Disabled. No Logical Drives Handled by BIOS.

Probable Cause: This warning message is displayed after you disable the **ROM** option in the configuration utility. When the **ROM** option is disabled, the BIOS cannot boot to Int 13h and cannot provide the ability to boot from the virtual disk. Int 13h is an interrupt signal that supports numerous commands that are sent to the BIOS and then passed to the physical disk. The commands include actions you can perform with a physical disk, such as reading, writing, and formatting.

Corrective Action: Enable the **ROM** option.

Configured disks removed or not accessible error message

Error Message: Some configured disks have been removed from your system or are no longer accessible. Check your cables and ensure all disks are present. Press any key or 'C' to continue.

Probable Cause: The message indicates that some configured disks were removed. If the disks were not removed, they are no longer accessible. The SAS cables for your system might be improperly connected.

Corrective Action: Check the cable connections and fix issues if any. Restart the system. If there are no cable problems, press any key or <C> to continue.

Dirty cache data error message

Error Message: The following virtual disks are missing: (x). If you proceed (or load the configuration utility), these virtual disks will be removed from your configuration. If you wish to use them at a later time, they will have to be imported. If you believe these virtual disks should be present, please power off your system and check your cables to ensure all disks are present. Press any key to continue, or 'C' to load the configuration utility. The cache contains dirty data, but some virtual disks are missing or will go offline, so the cached data cannot be written to disk. If this is an unexpected error, then please power off your system and check your cables to ensure all disks are present. If you continue, the data in cache will be permanently discarded. Press 'X' to acknowledge and permanently destroy the cached data.

Probable Cause: The controller preserves the dirty cache from a virtual disk if the disk becomes offline or is deleted because of missing physical disks. This message indicates that some configured disks were removed. If the disks were not removed, they are no longer accessible. The SAS cables for your system might be improperly connected.

Corrective Action: Check the cable connections and fix any problems. Restart the system. Use the HII configuration utility to import the virtual disk or discard the preserved cache. For the steps to discard the preserved cache, see [Clearing the cache memory](#).

Discovery error message

- 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 discovery did not complete within 120 seconds. The SAS cables for your system might be improperly connected.
- Corrective Action:** Check the cable connections and fix any problems. Restart the system.

Drive Configuration Changes Error Message

- Error Message:** Entering the configuration utility in this state will result in drive configuration changes. Press 'Y' to continue loading the configuration utility or please power off your system and check your cables to ensure all disks are present and reboot.
- Probable Cause:** The message is displayed after another HII warning indicating there are problems with previously configured disks and you have chosen to accept any changes and continue. The SAS cables for your system might be improperly connected.
- Corrective Action:** Check the cable connections and fix any problems before restarting the system. If there are no cable problems, press any key or <Y> to continue.

Windows operating system installation errors

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](#).

Firmware fault state error message

- Error Message:** Firmware is in Fault State.
- Corrective Action:** Contact [Global Technical Support](#).


Extra enclosure error message

- Error Message:** There are X enclosures connected to connector Y, but only maximum of 4 enclosures can be connected to a single SAS connector. Please remove the extra enclosures then restart your system.
- Probable Cause:** This message is displayed when the HII detects more than four enclosures connected to a single SAS connector.
- Corrective Action:** You must remove all additional enclosures and restart your system.

Foreign configuration found error message

- Error Message:** Foreign configuration(s) found on adapter. Press any key to continue, or 'C' to load the configuration utility or 'F' to import foreign configuration(s) and continue.
- Probable Cause:** When a controller firmware detects a physical disk with existing foreign metadata, it flags the physical disk as **foreign** and generates an alert indicating that a foreign disk was detected.
- Corrective Action:** Press **<F>** at this prompt to import the configuration (if all member disks of the virtual disk are present) without loading the **HII Configuration Utility**. Or press **<C>** to enter the **HII Configuration Utility** and either import or clear the foreign configuration.

Foreign configuration not found in HII error message

- 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 become inaccessible, the virtual disk become 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.

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 as they result in corrupted data and data loss. The following are the actions that occur in the case of multi-bit errors:

- 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 generates an event.
- If a multi-bit error occurs at run-time either in code/data or in the cache, the controller stops.
- The controller logs an event to the controller's internal event log and a message during POST is displayed indicating a multi-bit error has occurred.

 **NOTE:** In case of a multi-bit error, contact [Global 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 **pinned 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 [Clearing the cache memory](#).

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 passphrase 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 passphrase used to secure them. Supply the correct passphrase to import the secured foreign configuration. If you have lost or forgotten the passphrase, the secured foreign disks remain locked (inaccessible) until the appropriate passphrase is entered or if they are erased.
- **The secured virtual disk is in an offline state after supplying the correct passphrase**—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 [Creating virtual disks](#).

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 task 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 passphrase or unlocking a disk that had been previously locked.

Cryptographic Erase can be executed only on encryption-capable disks as long as the disks are not hot spares and are not configured or part of a virtual disk. Ensure that the conditions are met and see [Cryptographic Erase](#).

General issues

PERC card has yellow bang in device manager

Issue: The device is displayed in **Device Manager** but has a yellow bang (exclamation mark).
Corrective Action: Reinstall the driver. For more information on reinstalling drivers, see [Driver installation](#).

PERC card not seen in device manager

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

Physical disk issues

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 <https://www.dell.com/support> and replace the drive.

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.

Fatal error or data corruption reported

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

Physical disk displayed as blocked

Issue: One or more physical disks is displayed as **Blocked** and cannot be configured.
Corrective Action: Update the PERC cards to the latest firmware available on <https://www.dell.com/support>. Replace the drive and ensure not to use the blocked drive.

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.

Redundant virtual disk fails during reconstruction

Issue:	Multiple disks fails during a reconstruction process on a redundant virtual disk that has a hot spare.
Probable Cause:	Multiple physical disks in the virtual disk is failed or the cables are disconnected.
Corrective Action:	No action is required. The physical disk to which a reconstruction operation is targeted reverts to Ready state, and the virtual disk goes to Failed state. If there are any other virtual disks that can be supported by the capacity of the hot spare then the dedicated hot spare is converted to global hot spare, if not the hot spare will revert back to Ready state.

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 stress. There is only one rebuild I/O operation for every five host I/O operations.
Corrective Action:	If possible, reduce stress on the physical disk.

SMART errors


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

 **NOTE:** For information about where to find reports of SMART errors that could indicate hardware failure, see the Dell OpenManage storage management documentation at www.dell.com/openmanagemanuals.

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: <ol style="list-style-type: none">1. Back up your data.2. Force the physical disk offline.  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.  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 the topic Configuring hot spare.

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** or set up a global hot spare to replace the disk automatically.
 **NOTE:** For more information about the **Replace Member** feature, see the topic [Configuring hot spare](#).
 3. Replace the affected physical disk with a new physical disk of equal or higher capacity.
 4. Restore from the backup.

Replace member errors

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

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 message

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 Policies](#).

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 message

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 20. 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.

Table 17. 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 <i>i</i> 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
Blinks green slowly	The drive is rebuilding
Solid green	The drive is online
Blinks green for three seconds, amber for three seconds, and then powers off after six seconds	The rebuild has stopped

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.

System reports more drive slots than what is available

The system reports more slots than what is available in the following two scenarios:

- System drives are hot swappable with backplane.** When the system drives are hot swappable, the PERC controller is not able to communicate correctly with the backplane or enclosure. Hence, the PERC controller reports a generic enclosure with drive 16 slots. In iDRAC, under **Overview > Enclosures**, the **Enclosure ID** is displayed as **BP15G+0.0** and **Firmware version** is displayed as **03**.
- Corrective action** Turn off the system, reseal the controller and all the cables on the controller and backplane. If the issue is not resolved, contact your Dell Technical Service representative.

System drives are not hot swappable with cable direct attached.

When the system drives are not hot swappable, a default enclosure with 16 drive slots is expected to be reported (even though the system does not support that many drives).

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 10 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 18. 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

NOTE: The maximum number of virtual disks is currently limited to 192, because of the supported enclosure configuration.

RAID 10 configuration

In PERC 10 and PERC 11 controllers, RAID 10 can be configured without spanning up to 32 drives. 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 19. 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

Table 19. RAID 10 configurations (continued)

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
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 21. 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 22. 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 23. Example of Distributed Parity (RAID 5)

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




Figure 24. 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](#)
- [Receiving automated support with SupportAssist](#)

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 www.dell.com/recyclingworldwide 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:

1. Go to www.dell.com/support/home.
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 [Global 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 25. 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 EMC to route support calls to the appropriate personnel.

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

Receiving automated support with SupportAssist

Dell EMC SupportAssist is an optional Dell EMC Services offering that automates technical support for your Dell EMC server, storage, and networking devices. By installing and setting up a SupportAssist application in your IT environment, you can receive the following benefits:

- Automated issue detection — SupportAssist monitors your Dell EMC devices and automatically detects hardware issues, both proactively and predictively.
- Automated case creation — When an issue is detected, SupportAssist automatically opens a support case with Dell EMC Technical Support.
- Automated diagnostic collection — SupportAssist automatically collects system state information from your devices and uploads it securely to Dell EMC. This information is used by Dell EMC Technical Support to troubleshoot the issue.
- Proactive contact — A Dell EMC Technical Support agent contacts you about the support case and helps you resolve the issue.

The available benefits vary depending on the Dell EMC Service entitlement purchased for your device. For more information about SupportAssist, go to www.dell.com/supportassist.

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 EMC 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 20. Additional documentation resources for your system

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

Table 20. Additional documentation resources for your system (continued)

Task	Document	Location
	For information about installing the operating system, see the operating system documentation.	www.dell.com/operatingsystemmanuals
	For information about updating drivers and firmware, see the Methods to download firmware and drivers section in this document.	www.dell.com/support/drivers
Managing your system	For information about systems management software offered by Dell, see the Dell OpenManage Systems Management Overview Guide.	www.dell.com/poweredgemanuals
	For information about setting up, using, and troubleshooting OpenManage, see the Dell OpenManage Server Administrator User's Guide.	www.dell.com/openmanagemanuals > OpenManage Server Administrator
	For information about installing, using, and troubleshooting Dell OpenManage Enterprise, see the Dell OpenManage Enterprise User's Guide.	https://www.dell.com/openmanagemanuals
	For information about installing and using Dell SupportAssist, see the Dell EMC SupportAssist Enterprise User's Guide.	https://www.dell.com/serviceabilitytools
	For information about partner programs enterprise systems management, see the OpenManage Connections Enterprise Systems Management documents.	www.dell.com/openmanagemanuals
Working with the Dell PowerEdge RAID controllers	For information about understanding the features of the Dell PowerEdge RAID controllers (PERC), Software RAID controllers, or BOSS card and deploying the cards, see the Storage controller documentation.	www.dell.com/storagecontrollermanuals
Understanding event and error messages	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 .	www.dell.com/qrl
Troubleshooting your system	For information about identifying and troubleshooting the PowerEdge server issues, see the Server Troubleshooting Guide.	www.dell.com/poweredgemanuals