DELL Technologies

PowerEdge R760xs

Best choice in balanced compute and flexible storage for the most popular IT application

Buy the performance and flexibility you need

The new Dell PowerEdge R760xs is a 2U, two-socket rack server. Buy the best fit in scalable performance and large storage capability with this purpose-built 2U system. Focused on delivering the latest technology to power the most popular applications and workloads used by businesses today, including virtual desktop infrastructure (VDI), virtual machines (VMs), and software-defined storage (SDS). All delivered in a thoughtfully crafted platform that will provide balanced compute that fits in your current infrastructure.

Easily configurable

- Add up to two 5th generation Intel® Xeon® Scalable processors with up to 28 cores and 4th generation Intel® Xeon® Scalable processors with up to 32 cores per socket for faster performance
- Accelerate in-memory workloads with up to 16 DDR5 RDIMMS up to 5200 MT/sec
- Improve data throughput and reduce latency with support up to 8 I/O device (up to available 6 PCIe slots, 1 OCP 3.0 networking slot, and 1 dedicated PERC slot)
- Storage options include up to 12x 3.5” HDDs/SSDs, or up to 16x 2.5” HDD/SSDs, plus up to 8x NVMe drives

A breeze to cool

- Thoughtfully designed to fit in your current air-cooled infrastructure
- Alleviate the worry about expensive liquid cooling retrofitting to your data center
- Synchronize your workload needs with a tailored performance configuration that is air cooled
- Minimize the carbon footprint of your data center by better matching the system power consumption with anticipated workload requirements

Cyber Resilient Architecture for Zero Trust IT environment & operations

Security is integrated into every phase of the PowerEdge lifecycle, including protected supply chain and factory-to-site integrity assurance. Silicon-based root of trust anchors end-to-end boot resilience while Multi-Factor Authentication (MFA) and role-based access controls ensure trusted operations.

Increase efficiency and accelerate operations with autonomous collaboration

The Dell OpenManage™ systems management portfolio delivers a secure, efficient, and comprehensive solution for PowerEdge servers. Simplify, automate and centralize one-to-many management with the OpenManage Enterprise console and iDRAC.

Sustainability

From recycled materials in our products and packaging, to thoughtful, innovative options for energy efficiency, the PowerEdge portfolio is designed to make, deliver, and recycle products to help reduce the carbon footprint and lower your operation costs. We even make it easy to retire legacy systems responsibly with Dell Technologies Services.

Rest easier with Dell Technologies Services

Maximize your PowerEdge Servers with comprehensive services ranging from Consulting, to ProDeploy and ProSupport suites, Data Migration and more – available across 170 locations and backed by our 60K+employees and partners.

PowerEdge R760xs

The Dell PowerEdge R760xs offers compelling performance in a right-sized system with the latest PCIe Gen 5 bandwidth and large storage capability to support:

- Virtual Desktop Infrastructure (VDI)
- Virtual Machines (VMs)
- Software-Defined Storage Node

Feature	Technical Specifications	
Processor	Up to two 5th Generation Intel Xeon Scalable processor with up to 28 cores and 4th Generation Intel Xeon Scalable processor with up to 32 cores per processor	
Memory	• 16 DDR5 DIMM slots, supports RDIMM 1.5 TB max, speeds up to 5200 MT/s, supports registered ECC DDR5 DIMMs only	
Storage controllers	• Internal Controllers: PERC H965i, PERC H755, PERC H755N, PERC H355, HBA355i, HBA465i • Internal Boot: Boot Optimized Storage Subsystem (BOSS-N1): HWRAID 1, 2 x M.2 NVMe SSDs or USB • External HBA (non-RAID): HBA355e; Software RAID: S160	
GPU Options	2 x 75 W SW, LP	
Drive Bays	Front bays: • 0 drive bay • Up to 8 x 3.5-inch SAS/SATA (HDD/SSD) max 192 TB • Up to 12 x 3.5-inch SAS/SATA (HDD/SSD) max 288 TB • Up to 8 x 2.5-inch SAS/SATA/NVMe (HDD/SSD) max 122.88 TB • Up to 16 x 2.5-inch SAS/SATA (HDD/SSD) max 121.6 TB • Up to 16 x 2.5-inch (SAS/SATA) + 8 x 2.5-inch (NVMe) (HDD/SSD) max 244.48 TB	Rear bays: • Up to 2 x 2.5-inch SAS/SATA/NVMe (HDD/SSD) max 30.72 TB (supported only with 12 x 3.5-inch SAS/SATA HDD/SSD configuration)
Hot swap Redundant Power Supplies	• 1800 W Titanium 200—240 VAC or 240 VDC • 1400 W Titanium 100—240 VAC or 240 VDC • 1400 W Platinum 100—240 VAC or 240 VDC • 1400 W Titanium 277 VAC or HVDC (HVDC stands for High-Voltage DC, with 336V DC) • 1100 W Titanium 100—240 VAC or 240 VDC	• 1100 W -(48V — 60V) DC • 800 W Platinum 100—240 VAC or 240 VDC • 700 W Titanium 200—240 VAC or 240 VDC • 600 W Platinum 100—240 VAC or 240 VDC
Cooling Options	• Air cooling	
Fans	• Standard (STD) fans/High performance Silver (HPR) fans/ High performance Gold (VHP) fans, Up to 6 hot swappable fans	
Dimensions and Weight	• Height – 86.8 mm (3.41 inches) • Width – 482 mm (18.97 inches)	• Depth – 707.78 mm (27.85 inches) – without bezel 721.62 mm (28.4 inches) – with bezel • Weight – Max 28.6 kg (63.0 lbs.)
Form Factor	2U rack server	
Embedded Management	• iDRAC9 • iDRAC Direct • iDRAC RESTful API with Redfish	• iDRAC Service Module • Quick Sync 2 wireless module
Bezel	Optional LCD bezel or security bezel	
OpenManage Software	• CloudIQ for PowerEdge plug in • OpenManage Enterprise • OpenManage Enterprise Integration for VMware vCenter • OpenManage Integration for Microsoft System Center	• OpenManage Integration with Windows Admin Center • OpenManage Power Manager plugin • OpenManage Service plugin • OpenManage Update Manager plugin
Mobility	OpenManage Mobile	
OpenManage Integrations	• BMC Truesight • Microsoft System Center • OpenManage Integration with ServiceNow	• Red Hat Ansible Modules • Terraform Providers • VMware vCenter and vRealize Operations Manager
Security	• Cryptographically signed firmware • Data at Rest Encryption (SEDs with local or external key mgmt) • Secure Boot • Secure Erase	• Secured Component Verification (Hardware integrity check) • Silicon Root of Trust • System Lockdown (requires iDRAC9 Enterprise or Datacenter) • TPM 2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ
Embedded NIC	2 x 1 GbE LOM	
Network options	1 x OCP card 3.0 (optional)	
Ports	Front Ports: • 1 x iDRAC Direct (Micro-AB USB) port, 1 x USB 2.0, 1 x VGA Internal Ports: 1 x USB 3.0 (optional)	Rear Ports • 1 x Dedicated iDRAC Ethernet port, 1 x USB 2.0, 1 x USB 3.0, 1 x VGA, 1 x Serial (optional)
PCIe	• 1 CPU Configuration: Up to 4 PCIe slots (2 x8 Gen5, 1 x16 Gen4, 1 x8 Gen4) • 2 CPU configuration: Up to 6 PCIe slots (2 x16 Gen5, 3 x16 Gen4, 1 x8 Gen4)	
Operating System and Hypervisors	• Microsoft Windows Server with Hyper-V • Red Hat Enterprise Linux • SUSE Linux Enterprise Server	• VMware ESXi • Canonical Ubuntu Server LTS For specifications and interoperability details, see Dell.com/OSsupport.
OEM-ready version available	From bezel to BIOS to packaging, your servers can look and feel as if they were designed and built by you. For more information, visit Dell.com -> Solutions -> OEM Solutions.	

APEX Flex on Demand

Acquire the technology you need to support your changing business with payments that scale to match actual usage. For more information, visit https://www.delltechnologies.com/en-us/payment-solutions/flexible-consumption/flex-on-demand.htm.

Discover more about PowerEdge servers

Learn more about our PowerEdge servers

Learn more about our systems management solutions

Search our Resource Library

Follow PowerEdge servers on Twitter

Contact a Dell Technologies Expert for Sales or Support

June 2024

DELL EMC

THE INTEGRATED DELL REMOTE ACCESS CONTROLLER 9 (IDRAC9) WITH LIFECYCLE CONTROLLER

COMPLETE AGENT-FREE MANAGEMENT OF POWEREDGE SERVERS

Dell iDRAC9 provides security and intelligent automation.

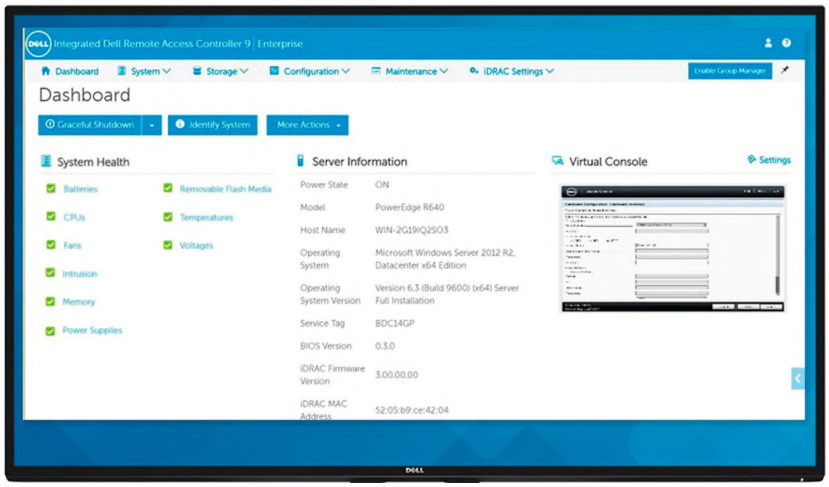

Modernize with the Dell EMC PowerEdge portfolio

The integrated Dell Remote Access Controller 9 (iDRAC9) delivers advanced, agent-free local and remote server administration. Embedded in every PowerEdge server, iDRAC9 provides a secure means to automate a multitude of common management tasks. Because iDRAC9 is embedded in every PowerEdge server, there's no additional software to install; just plug in power and network cables, and iDRAC9 is ready to go. Even before installing an operating system or hypervisor, IT administrators have a complete set of server management features at their fingertips: Maximize storage performance with up to 12 NVMe drives and ensure application performance scales easily.

- Configuration
- OS deployment
- Firmware updates
- Health monitoring
- Automation of other routine management activities

Scalable Architecture

With iDRAC9 in place across the PowerEdge portfolio, the same IT administration techniques and tools can be applied throughout This consistent management platform allows easy scaling of PowerEdge servers as an organization's infrastructure needs grow. Customers will be able to use the iDRAC RESTful API for the latest in scalable administration methods of PowerEdge servers. With this API, iDRAC9 enables support for the Redfish standard and enhances it with Dell EMC extensions to optimize at-scale management of PowerEdge servers. Regardless of size though, the entire OpenManage portfolio of systems management tools allows every customer to tailor an effective, affordable solution for their environment. This portfolio includes tools, consoles and integrations.

Each component leverages iDRAC9 to make management easy. By extending the reach of administrators to larger numbers of servers, that staff becomes more productive and drives down costs.

Intelligent Automation

Dell's agent-free management puts IT administrators in control. Once a PowerEdge server is connected to power and networking, that system can be monitored and fully managed, whether you're standing in front of the server or remotely over a network. In fact, with no need for software agents, an IT administrator can:

- Monitor
- Manage
- Update
- Troubleshoot and remediate Dell EMC servers

With features like zero-touch deployment and provisioning, iDRAC Group Manager, and System Lockdown, iDRAC9 is purpose-built to make server administration quick and easy. For those customers whose existing management platform utilizes in-band management, Dell EMC does provide iDRAC Service Module, a lightweight service that can interact with both iDRAC9 and the host operating system to support legacy management platforms.

Secure Local and Remote Management

Whether iDRAC9 is used via the updated, HTML5-based web interface, command line interface, or via a set of robust APIs such as the iDRAC RESTful API, security is ensured with HTTPS, SSL, Smart Card authentication, LDAP, and Active Directory integration. The iDRAC9 web interface, remote RACADM utility, and WS-MAN interfaces all support TLS 1.2. Every web page served by the iDRAC9 is delivered with TLS encryption at 256-bit strength (unless configured otherwise). Dell also supports encryption on the virtual KVM (virtual console redirection) and virtual media over TLS. The iDRAC9 virtual console and media also benefit from SSL encryption. Additionally, the iDRAC9 firmware is equipped with a default security certificate, which can be replaced by one of a customer's choosing. By providing secure access remote servers, administrators can carry out critical management functions while maintaining the integrity and security of their data.

The Heart of PowerEdge Manageability

The iDRAC9 provides common, embedded management across the PowerEdge family of servers, automation that lets your organization grow, and ensure security for peace of mind. This is why iDRAC is the core of managing Dell EMC servers. From the variety of tools and technologies in the OpenManage portfolio, a customer can build a management solution that matches their needs, and by leveraging iDRAC9, ensures optimal server management.

Key iDRAC9 Features and Specifications	
BIOS Recovery	Detect an invalid, untrusted BIOS image when a boot is attempted and recover to an authenticated, trusted BIOS image.
Connection View	Quickly check if server LOMs/NDCs and iDRAC are connected to the correct switches and ports via the GUI or by command line interface. This helps prevent costly remote dispatch of technicians to remediate cabling errors.
Full Power Cycle	By utilizing the iDRAC Service Module (iSM), DC power, including AUX power, can be temporarily removed via local or remote control to reset all power nodes in a server, saving time when troubleshooting.
iDRAC Direct	Secure front-panel USB connection to iDRAC web interface which eliminates the need for crash carts or a trip to the hot aisle of your data center. You can use the same port to insert a USB key to upload new system profile for secure, rapid system configuration
iDRAC Group Manager	Provides built-in, one-to-many monitoring and inventory of local iDRAC9s with no software to install. Ideal for customers who don't want to install and maintain a separate monitoring console. This feature does require iDRAC Enterprise licenses.
iDRAC RESTful API	With this API, iDRAC enables support for the Redfish standard and enhances it with Dell extensions.
Multi Vector Cooling	Airflow for each PCIe slot can be fine-tuned to ensure proper cooling. This allows for greater power efficiency and more precise cooling within each server for accessory cards.
OpenManage Mobile and Quick Sync 2	Use the OpenManage Mobile 2.0 (or higher) app on your handheld device to securely retrieve critical health data and easily perform bare-metal server configuration tasks via BLE/Wi-Fi connectivity. Compatible with various iOS and Android devices.
System Erase	With proper authentication, administrators can securely erase data from local storage (HDDs, SSDs, NVMs) and embedded flash devices.
System Lockdown	Helps to prevent configuration or firmware changes to a server when using Dell tools. Requires iDRAC Enterprise License.
Zero touch deployment and provisioning	When ordered with DHCP enabled from the factory, PowerEdge servers can be automatically configured when they are initially powered up and connected to your network. This process uses profile-based configurations that ensure each server is configured per your specifications. This feature requires an iDRAC Enterprise license.

iDRAC Licenses / Server Model	200-500 Series Rack / Tower	600+ Rack/Tower	Modular
Basic	Standard	n/a	n/a
Express	Optional	Standard	n/a
Express for Blades	n/a	n/a	Standard
Enterprise	Upgrade	Upgrade	Upgrade

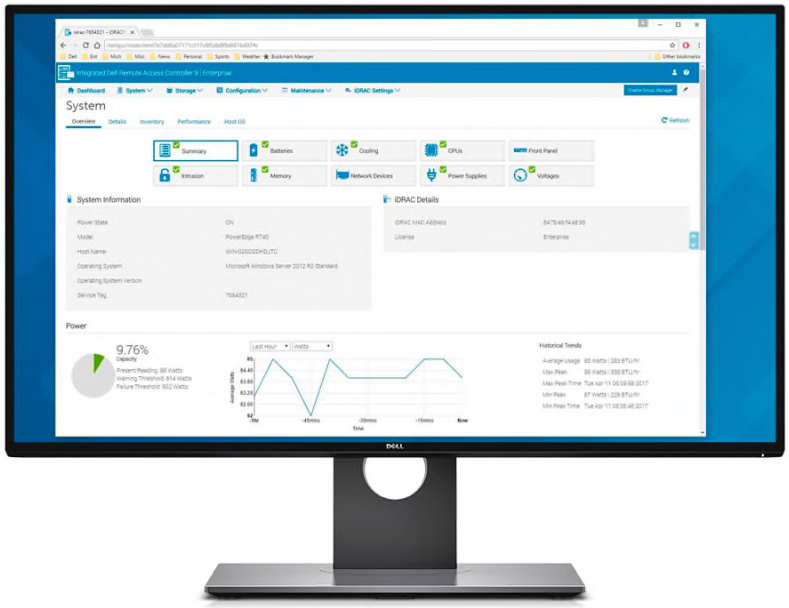

Learn more at
dell.com/poweredge and
delltechcenter.com/idrac

New features in yellow

iDRAC 9 License Levels and Features				
License Type	**Basic**	**Express**	**Express for Blades**	**Enterprise**
Interfaces / Standards				
Redfish	✓	✓	✓	✓
IPMI 2.0	✓	✓	✓	✓
DCMI 1.5	✓	✓	✓	✓
Web-based GUI	✓	✓	✓	✓
Racadm command line (local/remote)	✓	✓	✓	✓
SMASH-CLP (SSH-only)	✓	✓	✓	✓
Telnet	✓	✓	✓	✓
SSH	✓	✓	✓	✓
Serial Redirection	✓	✓	✓	✓
WSMAN	✓	✓	✓	✓
Network Time Protocol		✓	✓	✓
Connectivity				
Shared NIC	✓	✓	N/A	✓[1]
Dedicated NIC	✓	✓	✓	✓[2]
VLAN tagging	✓	✓	✓	✓
IPv4	✓	✓	✓	✓
IPv6	✓	✓	✓	✓
DHCP (new default; not static IP)	✓	✓	✓	✓
DHCP with Zero Touch				✓
Dynamic DNS	✓	✓	✓	✓
OS pass-through	✓	✓	✓	✓
iDRAC Direct - Front panel USB	✓	✓	✓	✓
Connection View	✓	✓		✓
NFS v4	✓	✓	✓	✓
SMB3.0 with NTLMv1 and NTLMv2	✓	✓	✓	✓
Security				
Role-based authority	✓	✓	✓	✓
Local users	✓	✓	✓	✓
SSL encryption	✓	✓	✓	✓
IP blocking		✓	✓	✓
Directory services (AD, LDAP)				✓
Two-factor authentication				✓
Single sign-on				✓
PK authentication		✓	✓	✓
Secure UEFI boot - certificate management	✓	✓	✓	✓
Lock down mode				✓
Unique iDRAC default password	✓	✓	✓	✓
FIPS 140-2	✓	✓	✓	✓
Customizable Security Policy Banner - login page	✓	✓	✓	✓

iDRAC 9 License Levels and Features				
License Type	**Basic**	**Express**	**Express for Blades**	**Enterprise**
Quick Sync 2.0 - optional auth for read operations	✓	✓	✓	✓
Quick Sync 2.0 - add mobile device number to LCL	✓	✓	✓	✓
System Erase of internal storage devices	✓	✓	✓	✓
Remote Presence				
Power control	✓	✓	✓	✓
Boot control	✓	✓	✓	✓
Serial-over-LAN	✓	✓	✓	✓
Virtual Media			✓	✓
Virtual Folders				✓
Remote File Share				✓
Virtual Console			✓	✓
HTML5 access to Virtual Console			✓	✓
VNC connection to OS				✓
Quality/bandwidth control				✓
Virtual Console collaboration (6 users)[2, 3]				✓
Virtual Console chat				✓
Virtual Flash partitions				✓
Group Manager				✓
HTTP / HTTPS support along with NFS/CIFS	✓	✓	✓	✓
Power & Thermal				
Real-time power meter	✓	✓	✓	✓
Power thresholds & alerts		✓	✓	✓
Real-time power graphing		✓	✓	✓
Historical power counters		✓	✓	✓
Power Capping				✓
OpenManage Power Center integration (view only)		✓	✓	✓
Temperature monitoring	✓	✓	✓	✓
Temperature graphing		✓	✓	✓
Health Monitoring				
Full agent-free monitoring	✓	✓	✓	✓
Predictive failure monitoring	✓	✓	✓	✓
SNMPv1, v2, and v3 (traps and gets)	✓	✓	✓	✓
Email Alerting		✓	✓	✓
Configurable thresholds	✓	✓	✓	✓
Fan monitoring	✓	✓	✓	✓
Power Supply monitoring	✓	✓	✓	✓
Memory monitoring	✓	✓	✓	✓
CPU monitoring	✓	✓	✓	✓
RAID monitoring	✓	✓	✓	✓
NIC monitoring	✓	✓	✓	✓
HD monitoring (enclosure)	✓	✓	✓	✓
Out of Band Performance Monitoring				✓
Alerts for excessive SSD wear	✓	✓	✓	✓

iDRAC 9 License Levels and Features				
License Type	**Basic**	**Express**	**Express for Blades**	**Enterprise**
Customizable settings for Exhaust Temperature	✓	✓	✓	✓
Update				
Remote agent-free update	✓	✓	✓	✓
Embedded update tools	✓	✓	✓	✓
Sync with repository (scheduled updates)				✓
Auto-update				✓
Improved PSU firmware updates	✓	✓	✓	✓
Deployment & Configuration				
Local configuration via F10	✓	✓	✓	✓
Embedded OS deployment tools	✓	✓	✓	✓
Embedded configuration tools	✓	✓	✓	✓
Auto-Discovery		✓	✓	✓
Remote OS deployment		✓	✓	✓
Embedded driver pack	✓	✓	✓	✓
Full configuration inventory	✓	✓	✓	✓
Inventory export	✓	✓	✓	✓
Remote configuration	✓	✓	✓	✓
Zerotouch configuration				✓
System Retire/Repurpose	✓	✓	✓	✓
Server Configuration Profile in GUI	✓	✓	✓	✓
Diagnostics, Service & Logging				
Embedded diagnostic tools	✓	✓	✓	✓
Part Replacement		✓	✓	✓
Server Configuration Backup				✓
Server Configuration Restore	✓	✓	✓	✓
Easy Restore (system configuration)	✓	✓	✓	✓
Easy Restore Auto Timeout	✓	✓	✓	✓
Health LED / LCD (requires optional bezel)[5]	✓	✓	N/A	✓
Quick Sync (require NFC bezel, 13G only)				
Quick Sync 2.0 (requires optional BLE/WiFi hardware)	✓	✓	✓	✓
iDRAC Direct (front USB management port)	✓	✓	✓	✓
iDRAC Service Module (iSM) embedded	✓	✓	✓	✓
Alert forwarding via iSM to inband monitoring consoles	✓	✓	✓	✓
Crash screen capture		✓	✓	✓
Crash video capture [4]				✓
Boot capture				✓
Manual reset for iDRAC (LCD ID button)	✓	✓	✓	✓
Remote reset for iDRAC (requires iSM)	✓	✓	✓	✓

iDRAC 9 License Levels and Features				
License Type	**Basic**	**Express**	**Express for Blades**	**Enterprise**
Virtual NMI	✓	✓	✓	✓
OS watchdog [4]	✓	✓	✓	✓
SupportAssist Report (embedded)	✓	✓	✓	✓
System Event Log	✓	✓	✓	✓
Lifecycle Log	✓	✓	✓	✓
Enhanced Logging in Lifecycle Controller Log	✓	✓	✓	✓
Work notes	✓	✓	✓	✓
Remote Syslog				✓
License management	✓	✓	✓	✓
Improved Customer Experience				
iDRAC -Faster processor, more memory	✓	✓	✓	✓
GUI rendered in HTML5	✓	✓	✓	✓
Add BIOS configuration to iDRAC GUI	✓	✓	✓	✓
iDRAC support for SW RAID licensing	✓	✓	✓	✓

footnotes:
1 Not available with blade servers.
2 500 series and lower rack and tower servers require a hardware card to enable this feature; this hardware offered at additional cost.
3 Requires vFlash SD card media.
4 Requires iDRAC Service Module (iSM) or OpenManage Server Administrator (OMSA).
5 Requires optional bezel.

Dell PowerEdge RAID Controller 11 User's Guide

PERC H755, H750, H355, and H350 Controller Series

Regulatory Model: UCPA-1101, UCPF-1100, UCPF-1110, UCPN-1100, UCSA-1111, UCSF-1100, and UCSM-1100
September 2023
Rev. A05

DELL Technologies

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.

Contents

Chapter 1: Dell Technologies PowerEdge RAID Controller 11 8
Features of PERC H755 adapter 9
Features of PERC H755 front SAS 9
Features of PERC H755N front NVMe 10
Features of PERC H755 MX adapter 10
Features of PERC H750 adapter SAS 11
Features of PERC H355 adapter SAS 11
Features of PERC H355 front SAS 12
Features of PERC H350 adapter SAS 12
Features of PERC H350 Mini Monolithic SAS 13
Operating systems supported by PERC 11 cards 13
Technical specifications of PERC 11 cards 14
Thermal specifications 16

Chapter 2: Applications and User Interfaces supported by PERC 11 18
Comprehensive Embedded Management 18
Dell OpenManage Storage Management 18
Human Interface Infrastructure Configuration Utility 18
The PERC Command Line Interface 19

Chapter 3: Features of PowerEdge RAID Controller 11 20
Controller features 20
Non-Volatile Memory Express 20
Opal Security Management 21
Hardware Root of Trust 21
1 MB I/O 21
Autoconfigure RAID 0 21
Disk roaming 22
FastPath 22
Non–RAID disks 23
Physical disk power management 23
Profile Management 23
Secure firmware update 23
Snapdump 23
Virtual disk features 23
Virtual disk write cache policy 24
Virtual disk read cache policy 24
Virtual disk migration 25
Virtual disk initialization 25
Full initialization 25
Fast initialization 25
Reconfigure virtual disks 26
Background operations 27
Background initialization 27

Consistency checks....28
Hard drive features....28
Self-Encrypting Disks....28
Instant secure erase....28
4 KB sector disk drives....28
Fault tolerance....29
The SMART feature....29
Patrol Read....29
Physical disk failure detection....30
Controller cache....31
Battery Transparent Learn Cycle....32
Linux operating system device enumeration....32

Chapter 4: Install and remove a PERC 11 card....34
Safety instructions....34
Before working inside your system....35
Remove the PERC H755 adapter....35
Install the PERC H755 adapter....36
Remove the PERC H755 front SAS card....37
Install the PERC H755 front SAS card....38
Remove the PERC H755N front NVMe card....39
Install the PERC H755N front NVMe card....41
Remove the PERC H755 MX adapter....42
Install the PERC H755 MX adapter....43
Remove the PERC H750 adapter SAS....45
Install the PERC H750 adapter SAS....45
Remove the PERC H355 adapter SAS....46
Install the PERC H355 adapter SAS....47
Remove the PERC H355 front SAS....48
Install the PERC H355 front SAS card....50
Remove the PERC H350 adapter SAS....51
Install the PERC H350 adapter SAS....52
Remove PERC H350 Mini Monolithic SAS....53
Install PERC H350 Mini Monolithic SAS....55

Chapter 5: Driver support for PERC 11....57
Creating the device driver media....57
Download and save PERC 11 drivers from the support site....57
Download and save PERC 11 drivers from the Dell Systems Service and Diagnostic Tools....57
Windows driver installation....58
Install PERC 11 driver while newly installing the Windows Server 2016 and later....58
Install PERC 11 driver on which the Windows Server 2016 is already installed and later....58
Update PERC 11 driver that runs on Windows Server 2016 or later....59
Linux driver installation....59
Install or update a RPM driver package using the KMOD support....60
Install or update a RPM driver package using the KMP support....60
Load the driver while installing an operating system....61

Chapter 6: Firmware....62

Upgrade firmware controller using Dell Update Package (DUP)....62

Chapter 7: Manage PERC 11 controllers using HII configuration utility....63
Enter the PERC 11 HII configuration utility....63
Exit the PERC 11 HII configuration utility....63
Navigate to Dell PERC 11 configuration utility....64
View the HII Configuration utility dashboard....64
Configuration management....65
Auto Configure RAID 0....65
Create virtual disks....65
Create profile based virtual disk....66
View disk group properties....67
Convert to Non–RAID disk....67
Delete configurations....67
Controller management....68
Clear controller events....68
Save controller events....68
Save debug log....68
Enable security....68
Disable security....68
Change security settings....69
Restore factory default settings....69
Auto configure behavior....69
Manage controller profile....69
Advanced controller properties....70
Virtual disk management....73
Virtual disk numbering....73
Configure Virtual Disks....75
Perform expand virtual disk operation....75
Perform consistency check....76
Physical disk management....76
View physical disk properties....76
Cryptographic erase....77
Physical disk erase....78
Assigning a global hot spare....78
Assigning a dedicated hot spare....78
Convert to Non–RAID disk....79
Hardware components....80
View battery properties....80
View physical disks associated with an enclosure....80
Security key management in HII configuration utility....81

Chapter 8: Security key and RAID management....82
Security key implementation....82
Local Key Management....82
Create a security key....82
Change Security Settings....83
Disable security key....83
Create a secured virtual disk....84

Secure a non-RAID disk....84
Secure a pre-existing virtual disk....84
Import a secured non-RAID disk....84
Import a secured virtual disk....85
Dell Technologies OpenManage Secure Enterprise Key Manager....85
Supported controllers for OpenManage Secure Enterprise Key Manager....85
Manage enterprise key manager mode....86
Disable enterprise key manager mode....86
Manage virtual disks in enterprise key manager mode....86
Manage non–RAID disks in enterprise key manager mode....86
Transition of drives from local key management to enterprise key management (without supported firmware for PERC and iDRAC)....86
Migrate of drives from local key management to enterprise key management (with supported firmware for PERC and iDRAC)....87

Chapter 9: Troubleshooting issues in PERC11 cards....88
Single virtual disk performance or latency in hypervisor configurations....88
Configured disks removed or not accessible error message....89
Dirty cache data error message....89
Discovery error message....89
Drive Configuration Changes Error Message....89
Windows operating system installation errors....90
Firmware fault state error message....90
Foreign configuration found error message....90
Foreign configuration not found in HII....90
Degraded state of virtual disks....90
Memory errors....91
Preserved Cache State....91
Security key errors....91
Secured foreign import errors....91
Failure to select or configure non Self-Encrypting Disks non-SED....91
Failure to delete security key....92
Failure of Cryptographic Erase on encryption-capable physical disks....92
General issues....92
PERC card has yellow bang in Windows operating system device manager....92
PERC card not seen in operating systems....92
Issues in controller, battery, and disk when operating at low temperature....92
Physical disk issues....92
Physical disk in failed state....92
Unable to rebuild a fault tolerant virtual disk....93
Fatal error or data corruption reported....93
Multiple disks are inaccessible....93
Rebuilding data for a failed physical disk....93
Virtual disk fails during rebuild using a global hot spare....94
Dedicated hot spare disk fails during rebuild....94
Redundant virtual disk fails during reconstruction....94
Virtual disk fails rebuild using a dedicated hot spare....94
Physical disk takes a long time to rebuild....94
Drive removal and insertion in the same slot generates a foreign configuration event....94
SMART errors....95

Smart error detected on a non–RAID disk....95
Smart error detected on a physical disk in a non–redundant virtual disk....95
Smart error detected on a physical disk in a redundant virtual disk....95
Replace member errors....96
Source disk fails during replace member operation....96
Target disk fails during replace member operation....96
A member disk failure is reported in the virtual disk which undergoes replace member operation....96
Linux operating system errors....96
Virtual disk policy is assumed as write-through....96
Unable to register SCSI device error message....97
Drive indicator codes....97
HII error messages....98
Unhealthy Status of the drivers....98
Rebuilding a drive during full initialization....98
System reports more drive slots than what is available....98
World Wide Number on drive sticker is not the same in applications....98
Backplane firmware revision not changing in PERC interfaces after an update....99

Chapter 10: Appendix RAID description....100
Summary of RAID levels....100
RAID 10 configuration....101
RAID terminology....102
Disk striping....102
Disk mirroring....102
Spanned RAID levels....103
Parity data....103

Chapter 11: Getting help....104
Recycling or End-of-Life service information....104
Contacting Dell....104
Locating the Express Service Code and Service Tag....104
Receiving automated support with SupportAssist....105

Chapter 12: Documentation resources....106

1

Dell Technologies PowerEdge RAID Controller 11

Dell Technologies PowerEdge RAID Controller 11, or PERC 11 is a series of RAID disk array controllers made by Dell for its PowerEdge servers. The PERC 11 series consists of the PERC H755 adapter, PERC H755 front SAS, PERC H755N front NVMe, PERC H750 adapter SAS, PERC H755 MX adapter, PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS cards that have the following characteristics:

- Provides reliability, high performance, and fault-tolerant disk subsystem management
- Offers RAID control capabilities including support for RAID levels 0, 1, 5, 6, 10, 50, 60
- Complies with Serial Attached SCSI (SAS) 3.0 providing up to 12 Gb/sec throughput
- Supports Dell-qualified Serial Attached SCSI (SAS), SATA hard drives, solid state drives (SSDs), and PCIe SSD (NVMe)
- Supports drive speeds of 8 GT/s and 16 GT/s at maximum x2 lane width for NVMe drives.

NOTE: Mixing disks of different speeds (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 (hard drive or SSD) is supported.

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

NOTE: PERC H750 adapter SAS, PERC H355 front SAS, PERC H355 adapter SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS do not support NVMe drives.

NOTE: RAID levels 5, 6, 50, and 60 are not supported on PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS.

NOTE: PERC H350 Mini Monolithic SAS has form factor variations (Low Profile) for specific platforms. For more information, see your platform manuals.

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

Topics:

- Features of PERC H755 adapter
- Features of PERC H755 front SAS
- Features of PERC H755N front NVMe
- Features of PERC H755 MX adapter
- Features of PERC H750 adapter SAS
- Features of PERC H355 adapter SAS
- Features of PERC H355 front SAS
- Features of PERC H350 adapter SAS
- Features of PERC H350 Mini Monolithic SAS
- Operating systems supported by PERC 11 cards
- Technical specifications of PERC 11 cards
- Thermal specifications

Features of PERC H755 adapter

This section describes the features of PERC H755 adapter.

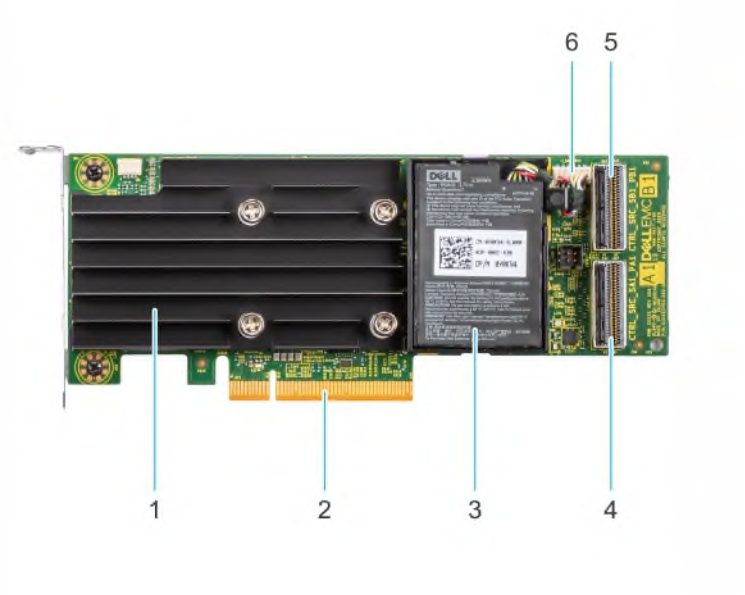

Figure 1. Features of PERC H755 adapter

1. Heatsink
2. PCIe connector
3. Battery
4. Backplane connector A
5. Backplane connector B
6. Battery cable connector

Features of PERC H755 front SAS

This section describes the features of PERC H755 front SAS.

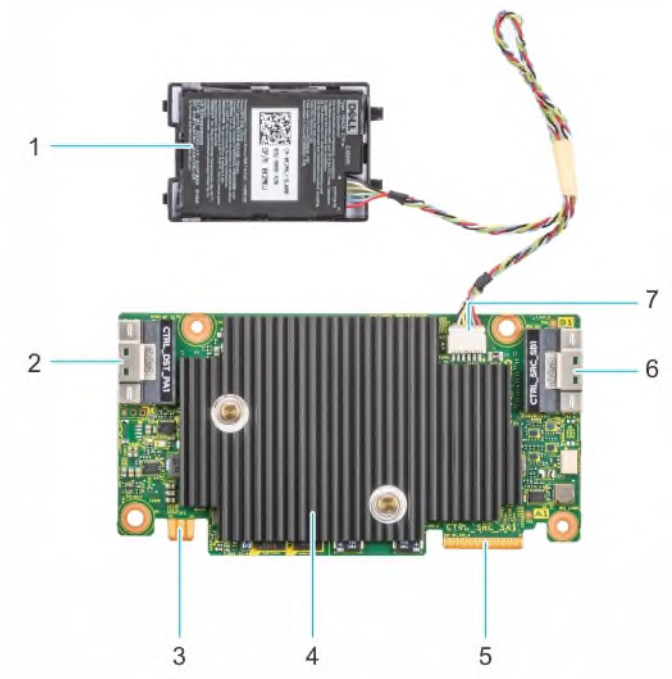

Figure 2. Features of PERC H755 front SAS

1. Battery
2. Backplane connector A

3. Power card edge connector
4. Heatsink
5. PCIe input connector
6. Backplane connector B
7. Battery cable connector

Features of PERC H755N front NVMe

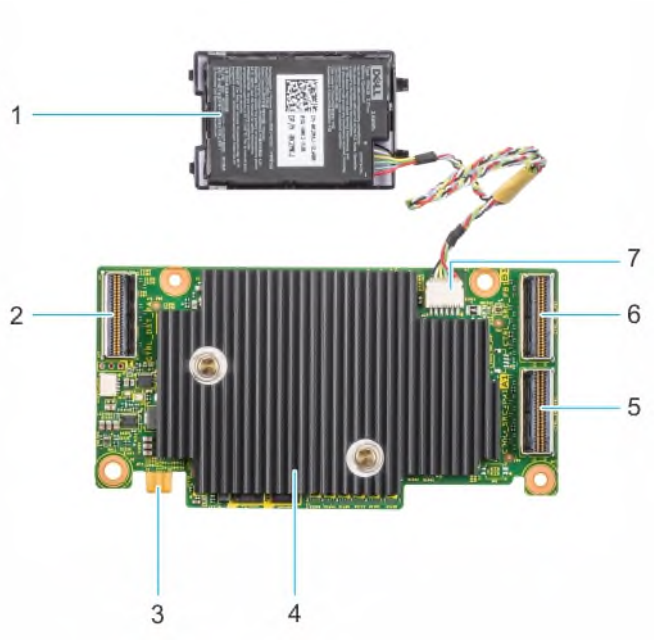

Figure 3. Features of PERC H755N front NVMe

1. Battery
2. PCIe cable connector
3. Power card edge connector
4. Heatsink
5. Backplane connector A
6. Backplane connector B
7. Battery cable connector

Features of PERC H755 MX adapter

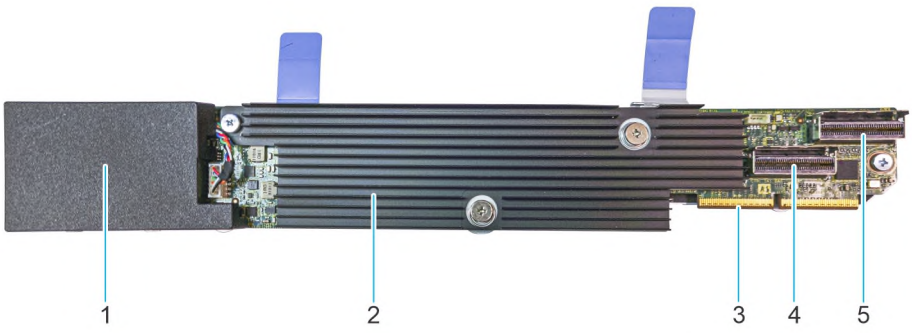

Figure 4. Features of PERC H755 MX adapter

1. Battery under cover
2. Heatsink
3. PCIe cable connector
4. Backplane connector A
5. Backplane connector B

Features of PERC H750 adapter SAS

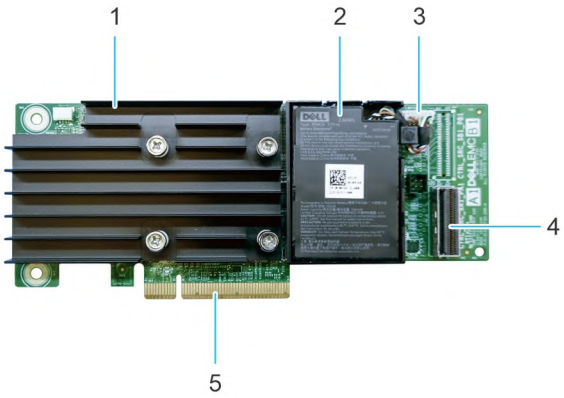

Figure 5. Features of PERC H750 adapter SAS

1. Heat sink
2. Battery
3. Battery cable connector
4. Backplane connector A
5. PCIe connector

Features of PERC H355 adapter SAS

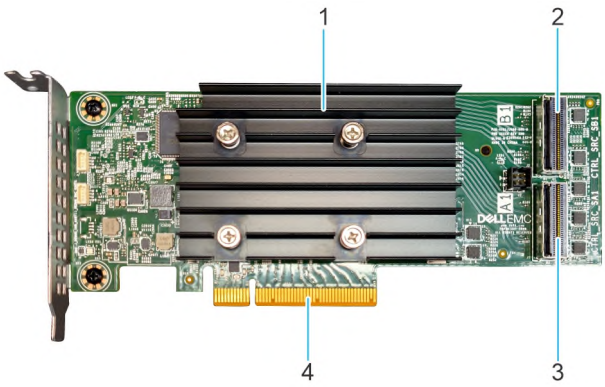

Figure 6. Features of PERC H355 adapter SAS

1. Heat sink
2. Backplane connector B
3. Backplane connector A
4. PCIe connector

Features of PERC H355 front SAS

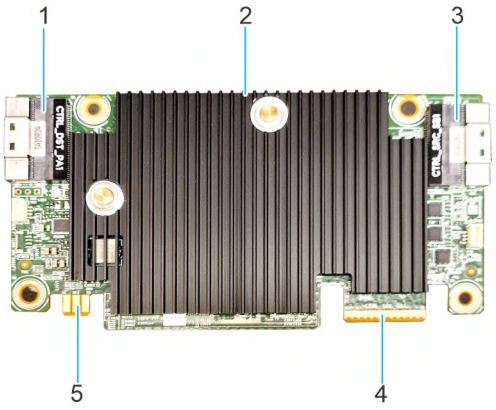

Figure 7. Features of H355 front SAS

1. PCIe input connector
2. Heat sink
3. Backplane connector B
4. Backplane connector A
5. Power card edge connector

Features of PERC H350 adapter SAS

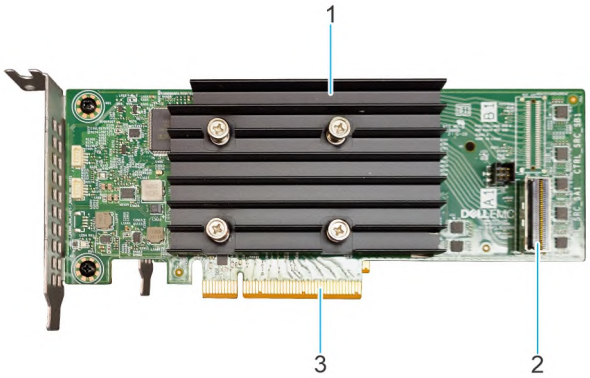

Figure 8. PERC H350 adapter SAS

1. Heat sink
2. Backplane connector A
3. PCIe connector

Features of PERC H350 Mini Monolithic SAS

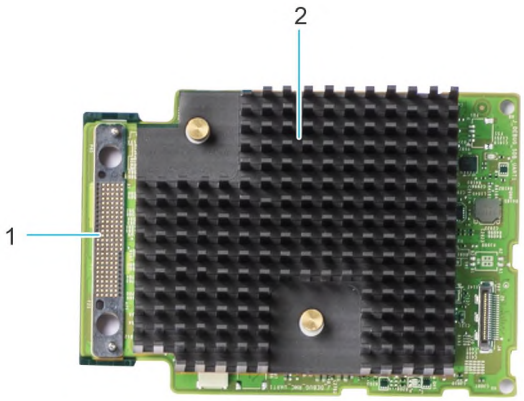

Figure 9. PERC H350 Mini Monolithic SAS

1. SAS cable connection
2. Heat sink

Operating systems supported by PERC 11 cards

See Dell Technologies Enterprise operating systems support for a list of supported operating systems by a specific server for the PERC 11 cards.

NOTE: For the latest list of supported operating systems and driver installation instructions, see the operating system documentation at Operating Systems Documentation. For specific operating system service pack requirements, see the Drivers and Downloads section on the support site.

Technical specifications of PERC 11 cards

The following table lists the specifications of PERC 11 cards:

Table 1. Technical specifications of PERC 11 cards

Feature	PERC H755 adapter	PERC H755 front SAS	PERC H755N front NVMe	PERC H755 MX adapter	PERC H750 adapter SAS
RAID levels	0, 1, 5, 6, 10, 50, 60	0, 1, 5, 6, 10, 50, 60	0, 1, 5, 6, 10, 50, 60	0, 1, 5, 6, 10, 50 ,60	0, 1, 5, 6, 10, 50 ,60
Non–RAID	Yes	Yes	Yes	Yes	Yes
Enclosures per port	Not applicable	Not applicable	Not applicable	Not applicable	Not applicable
Processor	Broadcom RAID-on-chip, SAS3916 chipset	Broadcom RAID-on-chip, SAS3916 chipset	Broadcom RAID-on-chip, SAS3916 chipset	Broadcom RAID-on-chip, SAS3916 chipset	Broadcom RAID-on-chip, SAS3916 chipset
Battery backup unit	Yes	Yes	Yes	Yes	Yes
Local Key Management security	Yes	Yes	Yes	Yes	Yes
Controller queue depth	5120	5120	5120	5120	5120
Secure enterprise key manager security	Yes	Yes	Yes	No	Yes
Non-volatile cache	Yes	Yes	Yes	Yes	Yes
Cache memory	8 GB DDR4 2666 MT/s cache	8 GB DDR4 2666 MT/s cache	8 GB DDR4 2666 MT/s cache	8 GB DDR4 2666 MT/s cache	8 GB DDR4 2666 MT/s cache
Cache function	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	Write back, write through, no read ahead, and read ahead
Max no of VDs in RAID mode	240	240	240	240	240
Max no of disk groups	240	240	240	240	240
Max no of VDs per disk group	16	16	16	16	16
Hot swap devices supported	Yes	Yes	Yes	Yes	Yes
Autoconfig	Yes	Yes	Yes	Yes	Yes
Hardware XOR engine	Yes	Yes	Yes	Yes	Yes
Online capacity expansion	Yes	Yes	Yes	Yes	Yes
Dedicated and global hot spare	Yes	Yes	Yes	Yes	Yes
Drives types	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS,	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS	Gen3 (8 GT/s) and Gen4 (16 GT/s) NVMe	3 Gbps SATA, 6 Gbps SATA/ SAS, and 12	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS

Table 1. Technical specifications of PERC 11 cards (continued)

Feature	PERC H755 adapter	PERC H755 front SAS	PERC H755N front NVMe	PERC H755 MX adapter	PERC H750 adapter SAS
	Gen3 (8 GT/s), and Gen4 (16 GT/s) NVMe			Gbps SAS, Gen3 (8 GT/s), and Gen4 (16 GT/s) NVMe	
VD strip size	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, and 1 MB	64 KB, 128 KB, 256 KB, 512 KB, 1 MB
PCIe support	Gen 4	Gen 4	Gen 4	Gen 4	Gen 4
SAS/SATA maximum drive support	• Without SAS Expander: 16 drives per controller • With SAS Expander: Limited by platform offerings	• Without SAS Expander: 16 drives per controller • With SAS Expander: Limited by platform offerings	Not applicable	Limited by platform: 8 drives per controller	• Without SAS Expander: 8 drives per controller • With SAS Expander: Limited by platform offerings
NVMe maximum drive support	• Without PCIe Switch Expander: 8 drives per controller • With PCIe Switch Expander: Limited by platform offerings	Not applicable	• Without PCIe Switch Expander: 8 drives per controller • With PCIe Switch Expander: Limited by platform offerings	Limited by platform:8 drives per controller	Not applicable

NOTE: PERC H755 adapter and PERC H755 MX supports either SAS, SATA, or NVMe drives depending on the backplane/server configuration.

NOTE: PERC controller supports only conventional magnetic recording (CMR) drives, and does not support shingled magnetic recording (SMR) drives.

NOTE: PERC H755 family of controllers currently support SEKM starting with firmware version 52.14.0-3901.

NOTE: For information about the number of drives in a disk group per virtual disk, see Summary of RAID levels

NOTE: As 14G PowerEdge Servers do not support Gen 4 speeds, PERC H750 adapter SAS will downtrain to Gen 3 speeds.

Table 2. Technical specifications of PERC 11 cards

Feature	PERC H355 adapter SAS	PERC H355 front SAS	PERC H350 adapter SAS	PERC H350 Mini Monolithic SAS
RAID levels	0, 1, 10	0, 1, 10	0, 1, 10	0, 1, 10
Non–RAID	Yes	Yes	Yes	Yes
Enclosures per port	Not applicable	Not applicable	Not applicable	Not applicable
Processor	Broadcom RAID-onchip, SAS3816 chipset	Broadcom RAID-onchip, SAS3816 chipset	Broadcom RAID-onchip, SAS3816 chipset	Broadcom RAID-onchip, SAS3816 chipset
Battery backup unit	No	No	No	No

Table 2. Technical specifications of PERC 11 cards (continued)

Feature	PERC H355 adapter SAS	PERC H355 front SAS	PERC H350 adapter SAS	PERC H350 Mini Monolithic SAS
Local Key Management security	No	No	No	No
Controller queue depth	1536	1536	1536	1536
Secure enterprise key manager security	No	No	No	No
Non-volatile cache	No	No	No	No
Cache memory	Not applicable	Not applicable	Not applicable	Not applicable
Cache function	Write through, no read ahead	Write through, no read ahead	write through, no read ahead	write through, no read ahead
Max no of VDs in RAID mode	32	32	32	32
Max no of disk groups	32	32	32	32
Max no of VDs per disk group	16	16	16	16
Hot swap devices supported	Yes	Yes	Yes	Yes
Autoconfig	Yes	Yes	Yes	Yes
Hardware XOR engine	Yes	Yes	Yes	Yes
Online capacity expansion	Yes	Yes	Yes	Yes
Dedicated and global hot spare	Yes	Yes	Yes	Yes
Drives types	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS, Gen3 (8 GT/s)	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS, Gen3 (8 GT/s)	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS, Gen3 (8 GT/s)	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS, Gen3 (8 GT/s)
VD strip size	64 KB	64 KB	64 KB	64 KB
PCIe support	Gen 4	Gen 4	Gen 4	Gen 4
SAS/SATA maximum drive support	• Without SAS Expander: 16 • With SAS Expander: Limited by platform offering	• Without SAS Expander: 16 • With SAS Expander: Limited by platform offering	• Without SAS Expander: 8 drives per controller • With SAS Expander: Limited by platform offering	• Without SAS Expander: 8 drives per controller • With SAS Expander: Limited by platform offering
NVMe maximum drive support	Not applicable	Not applicable	Not applicable	Not applicable

NOTE: As 14G PowerEdge Servers do not support Gen 4 speeds, PERC H350 adapter SAS and PERC H350 Mini Monolithic SAS will down train to Gen 3 speeds.

Thermal specifications

PERC 11 Controllers have an operating temperature range of 0C to 55C. System ambient temperatures may be less than or greater than these values.

NOTE: PERC Controllers may raise erroneous Battery, Disk, and Controller temperature errors if the controller is operating below the operational temperature range.

2

Applications and User Interfaces supported by PERC 11

PERC 11 card Management applications include the Comprehensive Embedded Management (CEM), Dell OpenManage Storage Management, The Human Interface Infrastructure (HII) configuration utility, and The PERC Command Line Interface (CLI). They enable you to manage and configure the RAID system, create and manage multiple disk groups, control and monitor multiple RAID systems, and provide online maintenance.

Topics:

- Comprehensive Embedded Management
- Dell OpenManage Storage Management
- Human Interface Infrastructure Configuration Utility
- The PERC Command Line Interface

Comprehensive Embedded Management

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

Using CEM enables you to do the following:

- Monitor devices with and without an operating systems 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 11 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 that you 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 OpenManage Manuals

Human Interface Infrastructure Configuration Utility

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

NOTE: The BIOS configuration utility <Ctrl> <R> is not supported on PERC 11 cards.

The PERC Command Line Interface

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

NOTE: For more information, see Dell PowerEdge RAID Controller CLI Reference Guide at Storage Controllers Manuals.

3

Features of PowerEdge RAID Controller 11

Topics:

- Controller features
- Virtual disk features
- Virtual disk initialization
- Reconfigure virtual disks
- Background operations
- Hard drive features
- Fault tolerance

Controller features

This section lists the following controller features supported on Dell Technologies PowerEdge RAID Controller 11 cards in detail:

- Non-Volatile Memory Express
- Opal Security Management
- Hardware Root of Trust
- 1 MB I/O
- Auto Configure RAID 0
- Disk roaming
- FastPath
- Non–RAID disks
- Physical disk power management
- Profile Management
- Secure firmware update
- Snapdump

Non-Volatile Memory Express

Non-Volatile Memory Express (NVMe) is a standardized, high-performance host controller interface and a storage protocol for communicating with non-volatile memory storage devices over the peripheral component interconnect express (PCIe) interface standard. The PERC 11 controller supports up to 8 direct-attach NVMe drives. The PERC 11 controller is a PCIe endpoint to the host, a PowerEdge server, and configured as a PCIe root complex for downstream PCIe NVMe devices connected to the controller.

NOTE: The NVMe drive on the PERC 11 controller shows up as a SCSI disk in the operating system, and the NVMe command line interface will not work for the attached NVMe drives.

Conditions under which a PERC supports an NVMe drive

- In NVMe devices the namespace identifier (NSID) with ID 1, which is (NSID=1) must be present.
- In NVMe devices with multiple namespace(s), you can use the drive capacity of the namespace with NSID=1.
- The namespace with NSID=1 must be formatted without protection information and cannot have the metadata enabled.
- PERC supports 512-bytes or 4 KB sector disk drives for NVMe devices.

Drive repair for NVMe initialization failure

If an NVME drive fails to initialize, the drive that is connected to PERC can be corrected in HII. The NVME initialization errors in the drives are listed as correctable and non-correctable errors in HII.

Repair drives with correctable NVMe initialization errors

Repair the drives with correctable NVMe initialization errors in HII to enable the drives to work properly.

About this task

Repairs can lead to permanent data loss in drives. Also, certain types of repairs can take a long time.

Steps

1. Log in to HII.
2. Go to **Main Menu** > **Hardware Components** > **Enclosure Management.**
   The drives with correctable and non-correctable errors are listed.
3. Select the drive and click **Repair**.
   If the repair is successful, the drive is listed under physical drives and removed from the correctable error list. If the drive has other correctable errors, the drive is listed again in the correctable errors list.
4. If the repair is not successful, click **Repair** again.

   NOTE: In case you want to stop the repair, stop the repair from the **Ongoing repairs** list.

   If the error is still not resolved or if the drive has other non-correctable errors, the drive is moved to the non-correctable error list.

Opal Security Management

Opal Security Management of Opal SED drives requires security key management support. You can use the application software or The Integrated Dell Remote Access Controller (iDRAC) to generate the security key that is set in the Opal drives and used as an authentication key to lock and unlock the Opal drives.

Hardware Root of Trust

Hardware RoT (RoT) builds a chain of trust by authenticating all the firmware components prior to its execution, and it permits only the authenticated firmware to perform and be flashed. The controller boots from an internal boot ROM (IBR) that establishes the initial root of trust and this process authenticates and builds a chain of trust with succeeding software using this root of trust.

1 MB I/O

PERC 11 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.

Autoconfigure RAID 0

The Autoconfigure RAID 0 feature creates a single drive RAID 0 on each hard drive that is in the ready state. For more information, see Auto Configure RAID 0.

NOTE: The Autoconfigure RAID 0 feature is not supported on PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS.

Autoconfigure behavior

The autoconfigure behavior automatically configures unconfigured drives during reboot and hot insertion. Unconfigured drives are configured according to the settings; but the configured drives remain unaffected. PERC 11 supports **Off and Non–RAID** settings.

Table 3. Autoconfigure behavior settings

Settings	Description
Off	Autoconfigure behavior is turned off.
Non–RAID	Unconfigured drives are configured as non–RAID disk during boot or during hot insertion; all the configured drives remain unaffected.
Off to Non–RAID disk	Unconfigured drives are converted to non–RAID disks; all the configured drives remain unaffected.
Non–RAID disk to Off	Unconfigured drives remain unconfigured good; all the configured drives remain unaffected.

NOTE: PERC H355 front SAS, PERC H355 adapter SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS converts an unconfigured good drive to non-RAID only if the drive has never been used before by that specific PERC.

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. If the physical disk is configured as a non–RAID disk, then the relocated physical disk is recognized as a non–RAID disk by the controller.

CAUTION: It is recommended that you perform 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

About this task

Perform the following steps to use disk roaming:

Steps

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.

Results

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 PERC 11 series of cards support FastPath.

To enable FastPath on a virtual disk, the cache policies of the RAID controller must be set to write-through and no 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.

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 11 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 OpenManage Manuals.

Profile Management

PERC 11 supports the PD240 and PD64 profiles. It defines controller queue depth and the maximum number of physical and virtual disks.

Table 4. Supported profile on PERC 11

Feature	PD240	PD64
Controller	PERC H755 front SAS, PERC H755 MX adapter, and PERC H750 adapter SAS	PERC H355 front SAS, PERC H355 adapter SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS
Maximum virtual disk supported	240	32
Controller queue depth	5120	1536

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 PowerEdge RAID Controller CLI Reference Guide at Storage Controllers Manuals.

Virtual disk features

This section lists the following virtual disk features supported on PERC 11 cards in detail:

- Virtual disk read cache policies
- Virtual disk write cache policies
- Virtual disk migration

- Virtual disk initialization
- Reconfiguration of virtual disks
- Background operations

Virtual disk write cache policy

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

Table 5. 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 policy

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

Table 6. 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.

NOTE: Adaptive read ahead is no longer supported. Selecting adaptive read ahead is equivalent to selecting the read ahead option.

Virtual disk migration

The PERC 11 series supports migration of virtual disks from one controller to another without taking the target controller offline. The controller can import RAID virtual disks in optimal, degraded, or partially degraded states. You cannot import a virtual disk that is offline. When a controller detects a configured physical disk, it marks the physical disk as foreign, and generates an alert indicating that a foreign disk was detected.

Disk migration pointers:

- Supports migration of virtual disks from H740P, H745, H745P MX, and H840 to the PERC 11 series except for H345.
- Supports migration of volumes that are created within the PERC 11 series.
- Does not support migration from the PERC 11 series to PERC H345, H740P, H745, H745P MX, and H840.
- Does not support migration from PERC H330, H730, and H830 to the PERC 11 series.

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

NOTE: Importing non–RAID drives and uneven span RAID 10 virtual disks from PERC 9 to PERC 11 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: Virtual disk migration from PERC H755 adapter, PERC H755 front SAS, PERC H755N front NVMe, PERC H750 adapter SAS, PERC H755 MX adapter to PERC H350 adapter SAS, PERC H350 Mini Monolithic SAS,PERC H355 front SAS, and PERC H355 adapter SAS 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 11 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 Configure 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 Configure 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 to create a virtual disk, the drive undergoes fast initialization. During this process all I/O requests to the drive will respond with a sense key of **"Not Ready"** and the I/O operation will fail. If the operating system attempts to read from the drive as soon as it discovers the drive, and while the fast initialization is still in process, then the I/O operation fails and the operating system reports an I/O error.

Reconfigure 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 following 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. Add additional physical disks to a virtual disk to increase its capacity.
3. After replacing all array members with larger drives than the original members, use the PERC CLI utility to expand the existing virtual disk to a larger size using the expandarray parameter. For more information, see Dell PowerEdge RAID Controller Command Line Interface Reference Guide.

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 0 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 7. 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 must be added.

Table 7. RAID level migration (continued)

Source RAID Level	Target RAID Level	Number of Physical Disks (Beginning)	Number of Physical Disks (End)	Capacity Expansion Possible	Description
RAID 0	RAID 6	1 or more	4 or more	Yes	Adds dual distributed parity redundancy; at least two disks must 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 must 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 performed.

NOTE:

- 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.

- PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS background operations will not run until the operating system boots.

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 Perform 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 11 cards in detail:

- Self-Encrypting Disks (SED)
- Instant Secure Erase (ISE)
- 4 KB sector disk drives

Self-Encrypting Disks

Select PERC 11 cards support self-encrypting disks (SEDs) for protection of data against loss or theft of SEDs. For information about cards supported, see Technical specifications. Protection is achieved by the use of encryption technology on the drives. There is one security key per controller. You can manage the security key using local key management (LKM) or OpenManage Secure Enterprise Key Manager also referred as Secure Enterprise Key Manager (SEKM). The controller use the security key to lock and unlock access to encryption-capable physical disks. To take advantage of this feature, you must:

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

PERC cannot use SEDs that are secured by a non-PERC entity. Ensure that the SED is reprovisioned in an applicable manner by that non-PERC entity before connecting to PERC.

For more information, see the Security key and RAID management section.

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

NOTE: PERC 11 supports Trusted Computing Group Enterprise (TCG) Security Subsystem Classes (SSC) SAS or SATA SED drives and TCG Opal SSC NVMe 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 11 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, see 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 under the select boot device option. For more information, see Enable boot support.

Fault tolerance

The PERC 11 series supports the following:

- Self-Monitoring and Reporting Technology (SMART)
- Patrol read
- Physical disk failure detection
- Controller cache
- Battery Transparent Learn Cycle

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 hard drive failures. Data on SMART compliant hard drives 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 that are related to read/write surface failure, such as seek error rate and excessive bad sectors.

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

NOTE: PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS controllers do not monitor predictive failures for non-RAID disks.

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 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 storage management applications 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 OpenManage Manuals.

Physical disk failure detection

If a disk fails and it is replaced with a new disk, the controller will automatically start a rebuild on the new disk. See, Configured slot behavior. Automatic rebuilds can also 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 11 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 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 OpenManage Manuals.

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 8. 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 • Original drive data lost
Insert configured locked drive into the system (unlockable)	Foreign	Cryptographic Erase (If configured VD is not secured) • Rebuild or copyback start • Original drive data lost
Insert locked drive into the system (non-unlockable)	Foreign locked	Foreign locked

Physical disk hot swapping

Hot swapping is the manual replacement of a disk while the PERC 11 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 11 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 NVMe drive can replace a NVMe drive.

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 PERC 11 series of cards contain local DRAM on the controllers. This DRAM can cache I/O operations 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:

- 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. This action degrades the performance. If a foreign disk is present, it is recommended that you import, clear, or remove the foreign disk to prevent an impact on the performance.

Controller cache preservation

The controller is capable of preserving its cache in the event of a system power outage or improper system shutdown. The PERC 11 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

About this task

Complete these steps if a system power loss or improper system shutdown has occurred.

Steps

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 about 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.

Linux 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 ID.

Enumeration order for PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS

Steps

1. Non-RAID disks are enumerated first.
2. Virtual disks (VDs) are enumerated second, based on virtual disk target ID. Target IDs are assigned to the VDs in the ascending order when they are created. The first created VD is assigned the lowest available target ID, and the last created VD is assigned the highest available target ID. The first created VD is discovered first by the operating system.

   NOTE: The PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS non-RAID disks may not appear in the slot order.

Enumeration order for PERC H755 front SAS, PERC H755 front SAS, PERC H755N front NVMe, PERC H750 adapter SAS, and PERC H755 MX adapter

This section describes the order of enumerating PERC H-series controlers.

Steps

1. Non-RAID disks are enumerated first based on slot ID.
2. Virtual disks (VDs) are enumerated, second based on the virtual disk target ID. Target IDs are assigned to the VDs in the descending order when they are created. The first created VD is assigned the highest available target ID, and the last created VD is assigned the lowest available target ID. Therefore, the last created VD is discovered first by the operating system.

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

4

Install and remove a PERC 11 card

Topics:

- Safety instructions
- Before working inside your system
- Remove the PERC H755 adapter
- Install the PERC H755 adapter
- Remove the PERC H755 front SAS card
- Install the PERC H755 front SAS card
- Remove the PERC H755N front NVMe card
- Install the PERC H755N front NVMe card
- Remove the PERC H755 MX adapter
- Install the PERC H755 MX adapter
- Remove the PERC H750 adapter SAS
- Install the PERC H750 adapter SAS
- Remove the PERC H355 adapter SAS
- Install the PERC H355 adapter SAS
- Remove the PERC H355 front SAS
- Install the PERC H355 front SAS card
- Remove the PERC H350 adapter SAS
- Install the PERC H350 adapter SAS
- Remove PERC H350 Mini Monolithic SAS
- Install PERC H350 Mini Monolithic SAS

Safety instructions

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

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

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

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

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

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

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

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

Before working inside your system

Prerequisites

Follows the steps listed in Safety instructions.

Steps

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

Remove the PERC H755 adapter

Prerequisites

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

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

Steps

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

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

2. Open the system.
3. Locate the PERC card in the 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 cables connected to the card:
   a. Press down and hold the metal tab on the cable connector.
   b. Pull the cable out of the connector.
6. Replace the storage controller card and reconnect the data cables before placing them in the riser. For more information on installing the card, see Install PERC H755 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.

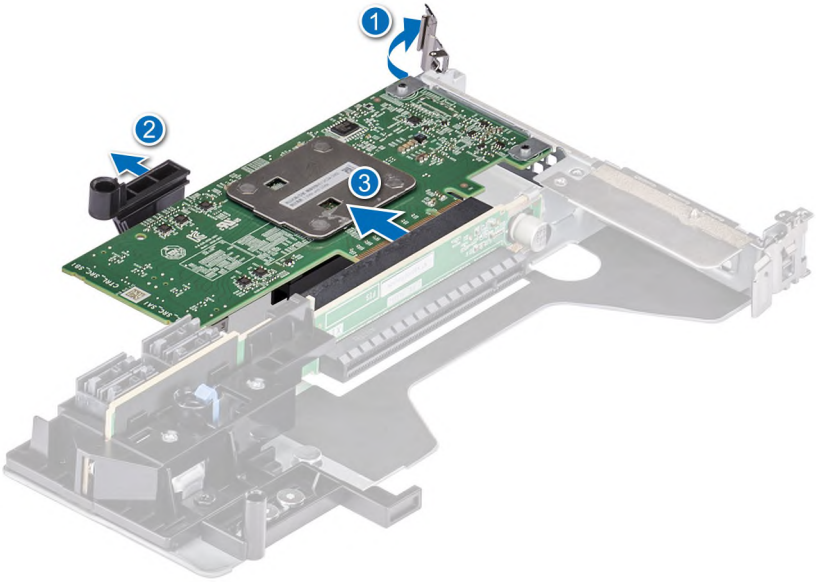

Figure 10. Remove the PERC H755 adapter

Install the PERC H755 adapter

Prerequisites

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

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

Steps

1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.
2. Open the system.
3. Align the card-edge connector with the connector on the 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 data cable connectors to the card.
6. Route the data cable through the channel on the inner side of the chassis to the backplane.
7. Attach the connector to the corresponding connector on the backplane as labeled on the controller.
8. Close the system.
9. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.

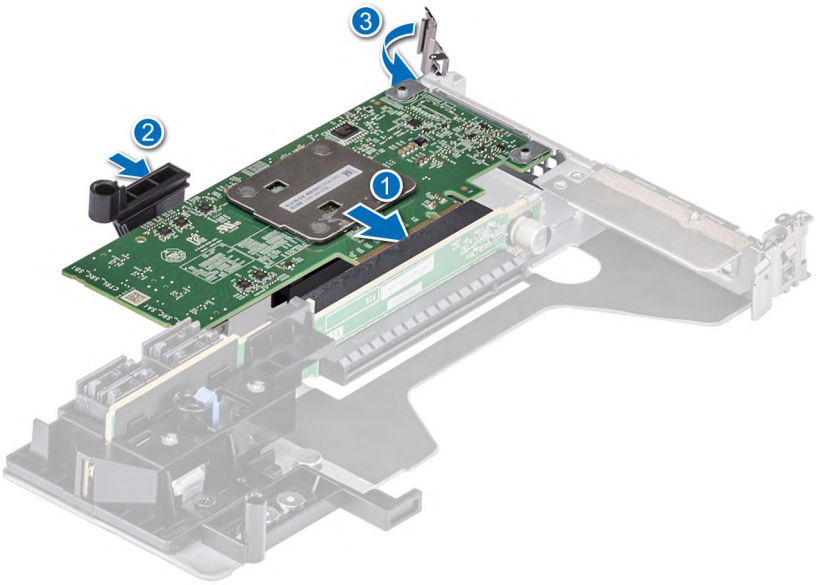

Figure 11. Install the PERC H755 adapter

Remove the PERC H755 front SAS card

Prerequisites

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

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

Steps

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

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

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

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

4. Unscrew the fasteners on the controller carrier and slide the carrier away from the backplane, disconnecting the controller from the backplane.

   If you are removing a PERC H755 front SAS controller in the upside down orientation, you must remove both the backplane and the controller at the same time because of the limited clearance available:

   a. Uninstall all drives from the backplane.
   b. Disconnect all cables between the PERC and the backplane.
   c. Lift the backplane and PERC from the system.

5. Disconnect any cables connected to the card:
   a. Press down and hold the metal tab on the cable connector.
   b. Pull the cables out of the connector.
6. Remove the PERC controller from the controller carrier.

7. Insert the replacement controller into the carrier and secure it with the appropriate screws.
8. Take the replacement storage controller and reconnect the cables before reconnecting it to the backplane.

   If you are removing a PERC H755 front SAS controller in the 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 Install PERC H755 front SAS card.
9. Close the system.
10. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.

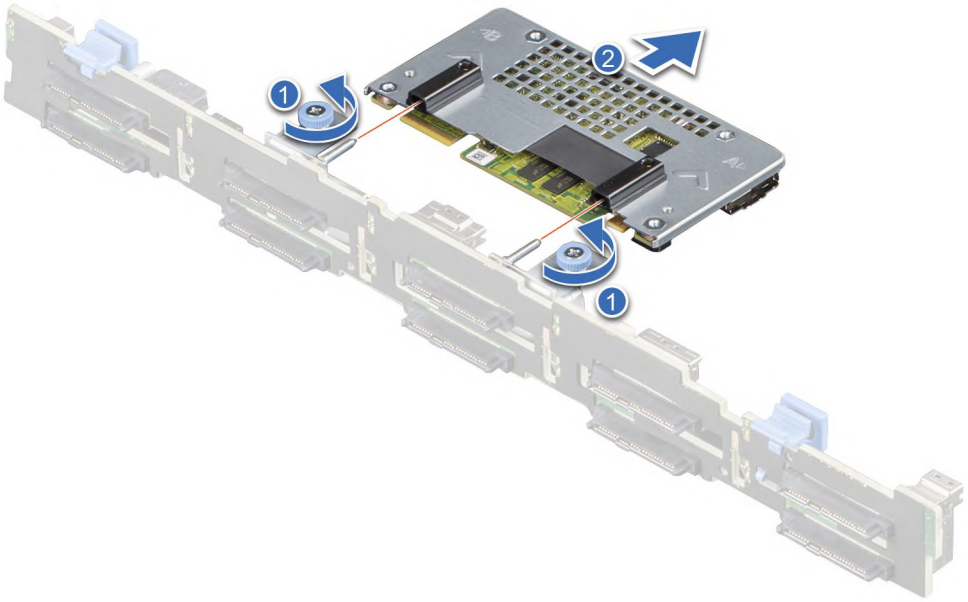

Figure 12. Remove the PERC H755 front SAS card

Install the PERC H755 front SAS card

Prerequisites

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

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

Steps

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

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

2. Open the system.
3. Connect the PERC card to the carrier and ensure that the screws are properly fastened in place.

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

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

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

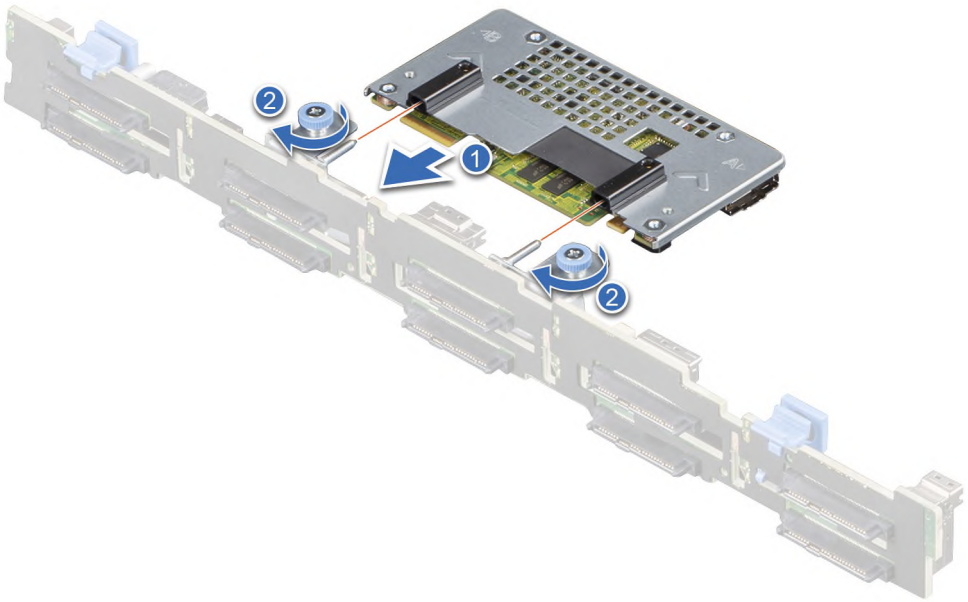

Figure 13. Install the PERC H755 front SAS card

Remove the PERC H755N front NVMe card

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or

telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

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

Steps

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

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

2. Open the 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 to disconnect the controller from the backplane.

   If you are removing a PERC H755N front NVMe controller in the upside down orientation, you must remove both the backplane and the controller at the same time because of the limited clearance available:

   a. Uninstall all drives from the backplane.
   b. Disconnect all cables between the PERC and the backplane.
   c. Lift the backplane and PERC from the system.

5. Disconnect any cables connected to the card:
   a. Press down and hold the metal tab on the cable connector.
   b. Pull the cable out of the connector.
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 cable before reconnecting it to the backplane.

   If you are removing a PERC H755 front NVMe controller in the upside down orientation, reattach the PERC controller to the backplane first before reinstalling the backplane into the system.

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

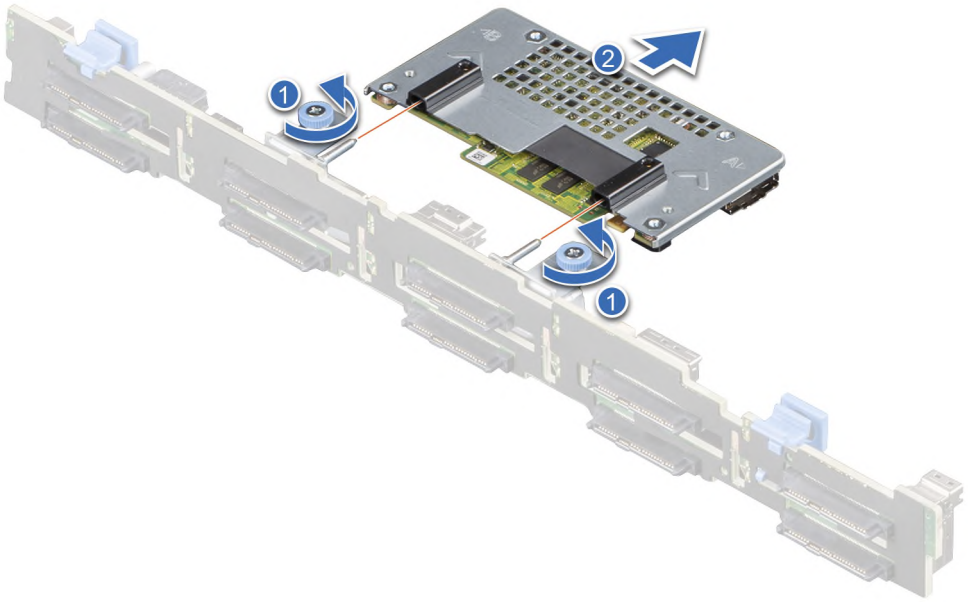

Figure 14. Remove the PERC H755N front NVMe card

Install the PERC H755N front NVMe card

Prerequisites

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

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

Steps

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

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

2. Open the system.
3. Connect the PERC card to the carrier and ensure 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 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.

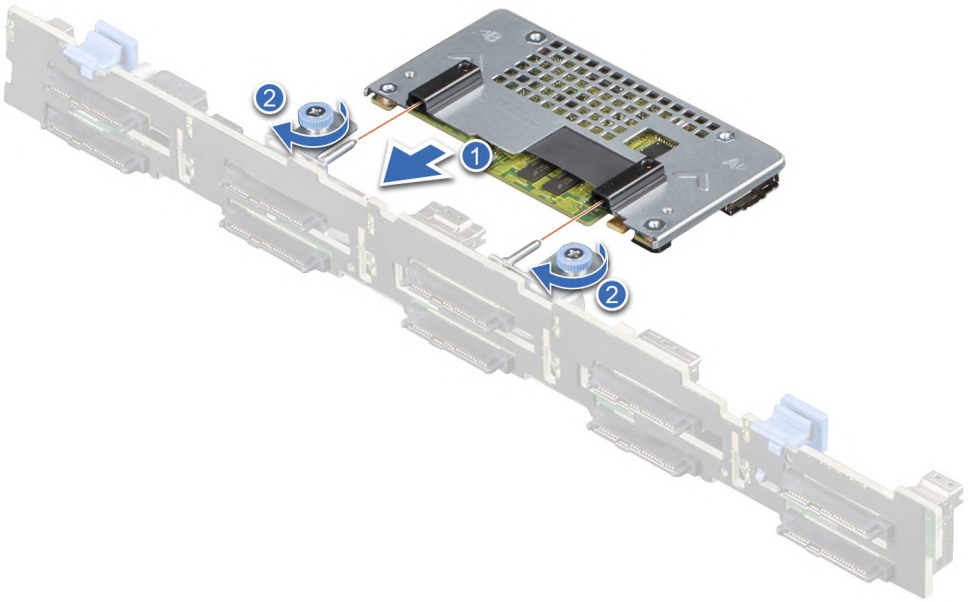

Figure 15. Install the PERC H755N front NVMe card

Remove the PERC H755 MX adapter

Prerequisites

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

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

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

Steps

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

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

2. Open the sled.
3. Locate the PERC card on the system board.

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

4. Using the blue tab, rotate the lever of the controller.
5. Pull the release lever upward to disengage the controller from the connector.
6. Disconnect the cable from the card. To disconnect the cable:
   a. Press and hold the metal tab on the cable connector.
   b. Pull the cable out of the connector.
7. Lift the card from the system board.
8. Replace the storage controller card and connect the cable. For information on installing the card, see Install the PERC H755 MX adapter.
9. Close the sled.
10. Insert the sled into the MX chassis and turn on the system and any attached MX chassis peripherals.

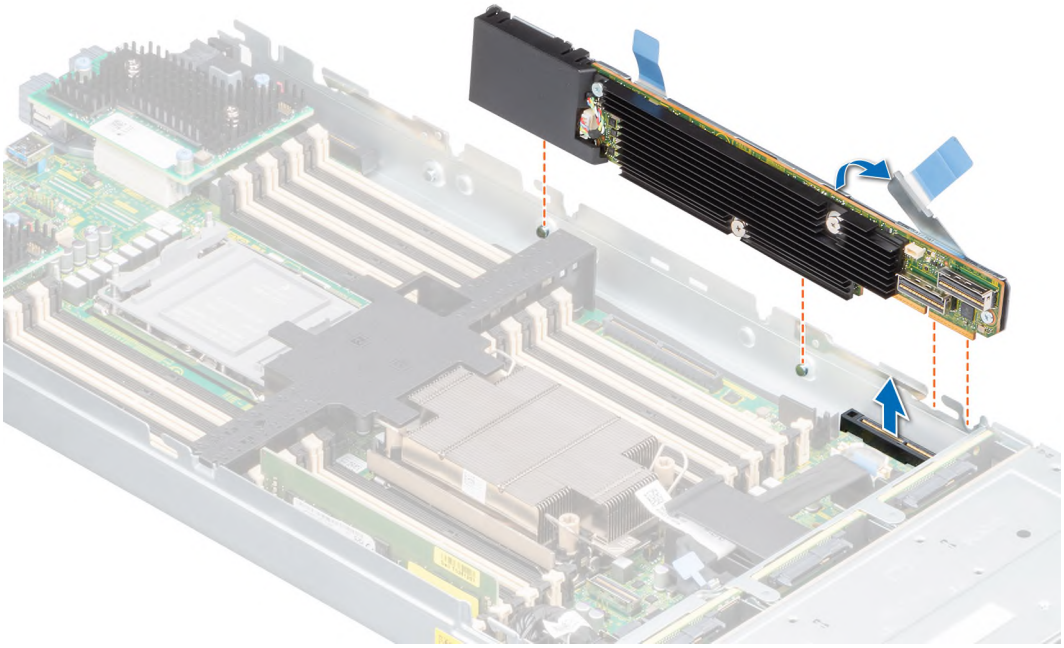

Figure 16. Remove the PERC H755 MX adapter

Install the PERC H755 MX adapter

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or

telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

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

Steps

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

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

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

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

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

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

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

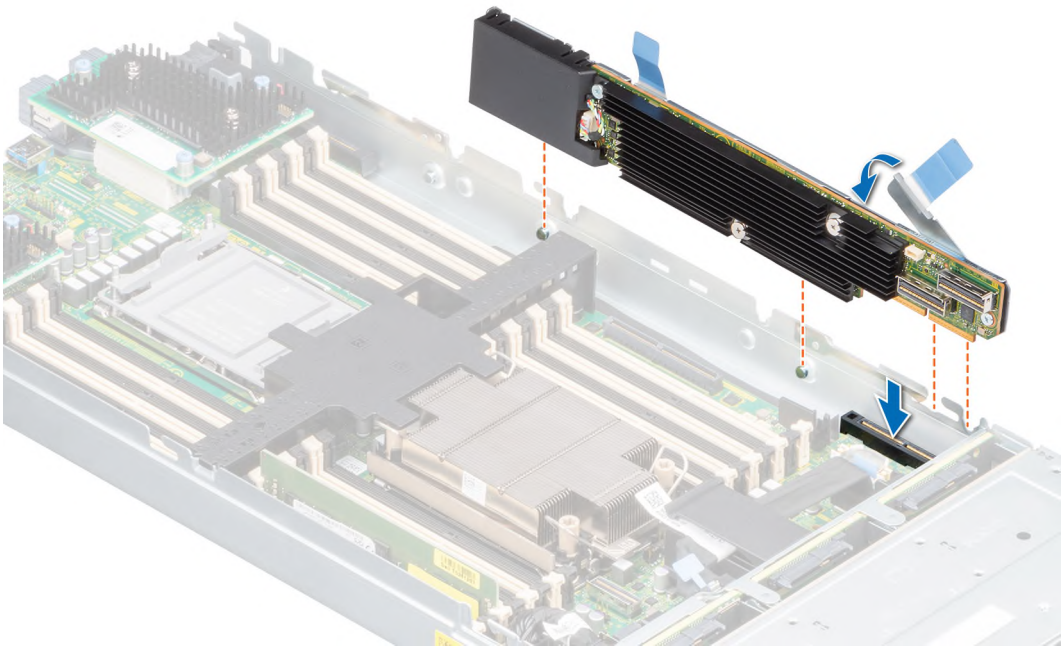

Figure 17. Install the PERC H755 MX adapter

Remove the PERC H750 adapter SAS

Prerequisites

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

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

Steps

1. 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 Install the H750 adapter SAS.
7. Close the system.
8. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.

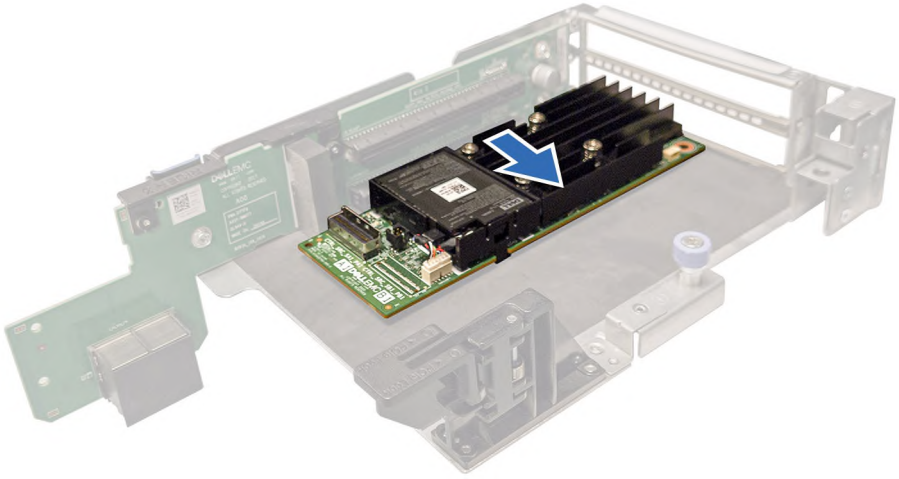

Figure 18. Remove PERC H750 adapter SAS

Install the PERC H750 adapter SAS

Prerequisites

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

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

Steps

1. 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 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.

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

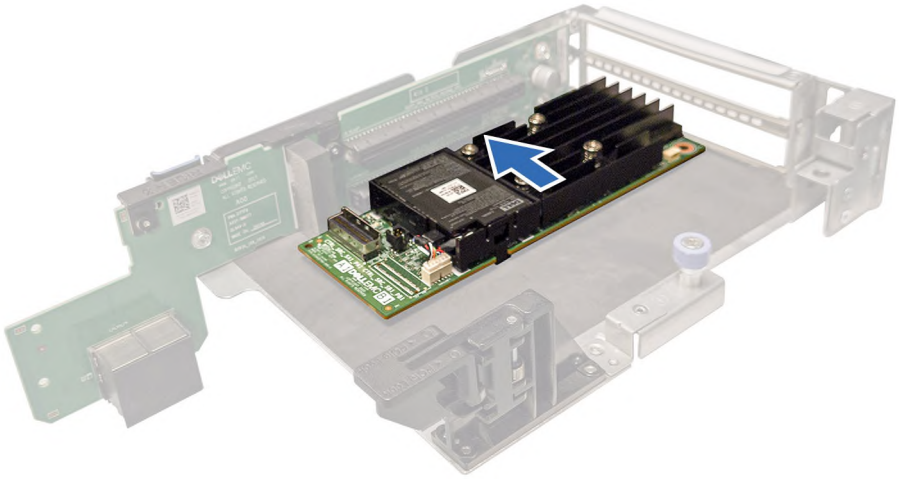

Figure 19. Install PERC H750 adapter SAS

Remove the PERC H355 adapter SAS

Describes the tasks to remove a PERC H355 adapter SAS controller from a server.

Prerequisites

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

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

Steps

1. 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 Install the PERC H355 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.

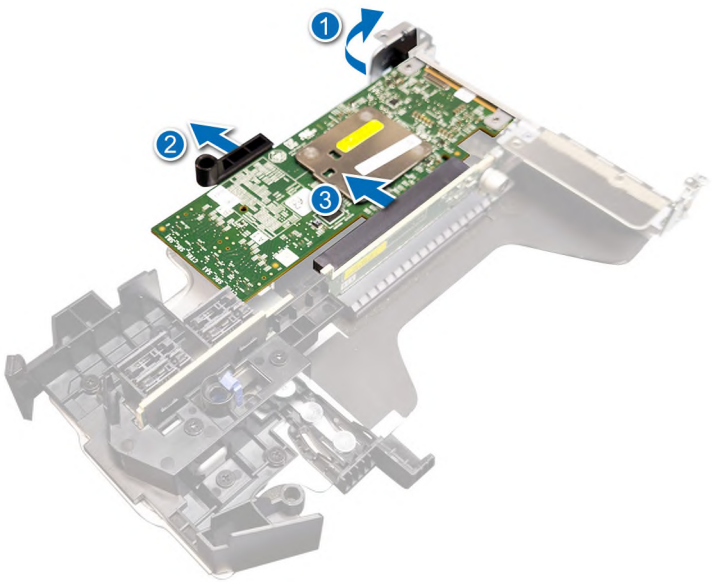

Figure 20. Remove the PERC H355 adapter SAS

Install the PERC H355 adapter SAS

Describes the tasks to install a PERC H355 adapter SAS controller in a server.

Prerequisites

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

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

Steps

1. 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.

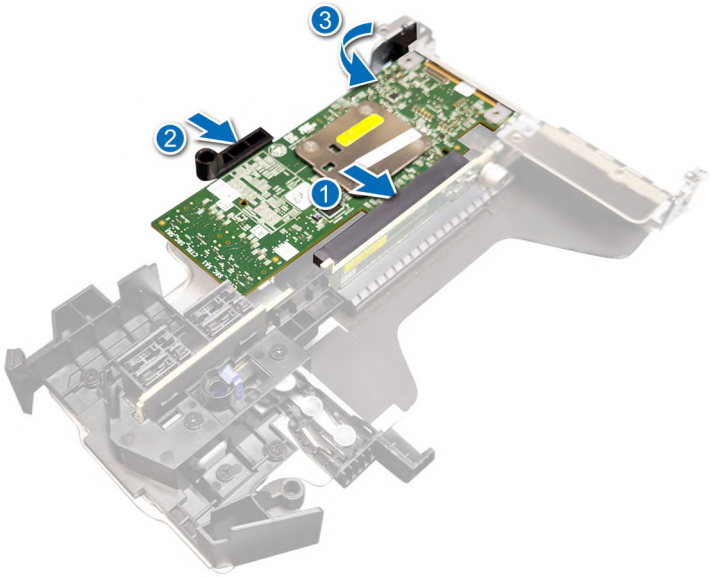

Figure 21. Install the PERC H355 adapter SAS

Remove the PERC H355 front SAS

Prerequisites

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

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

Steps

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

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

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

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

4. Unscrew the fasteners on the controller carrier and slide the carrier away from the backplane, disconnecting the controller from the backplane.

   If you are removing a PERC H355 front SAS controller in the upside down orientation, you must remove both the backplane and the controller at the same time because of the limited clearance available:

   a. Uninstall all drives from the backplane.
   b. Disconnect all cables between the PERC and the backplane.
   c. Lift the backplane and PERC from the system.

5. Disconnect any cables connected to the card:
   a. Press down and hold the metal tab on the cable connector.
   b. Pull the cables out of the connector.
6. Remove the PERC controller from the controller carrier.
7. Insert the replacement controller into the carrier and secure it with the appropriate screws.
8. Take the replacement storage controller and reconnect the cables before reconnecting it to the backplane.

   If you are removing a PERC H355 front SAS controller in the 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 Install the PERC H355 front.
9. Close the system.
10. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.

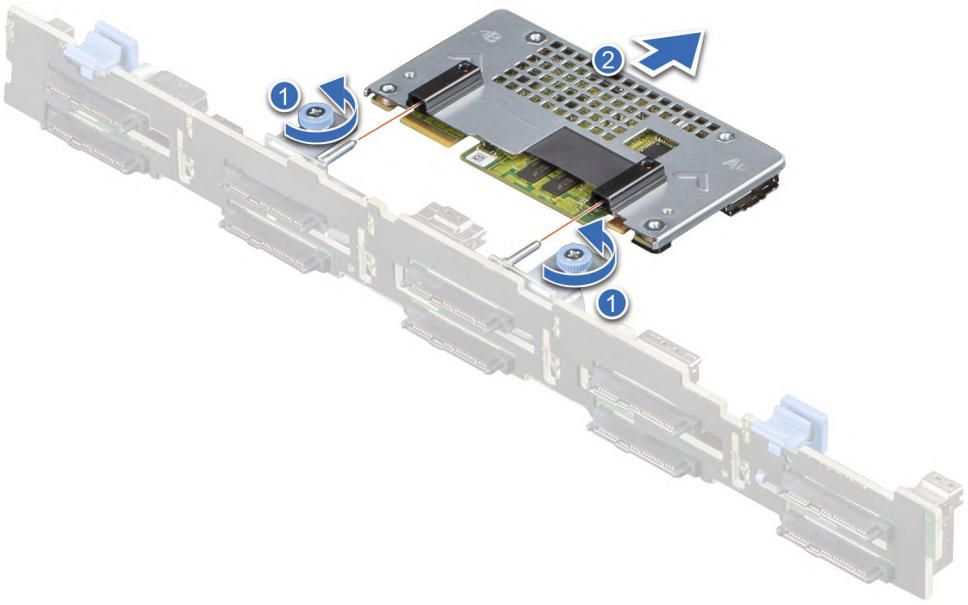

Figure 22. Remove the PERC H355 front SAS

Install the PERC H355 front SAS card

Prerequisites

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

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

Steps

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

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

2. Open the system.
3. Connect the PERC card to the carrier and ensure that the screws are properly fastened in place.

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

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

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

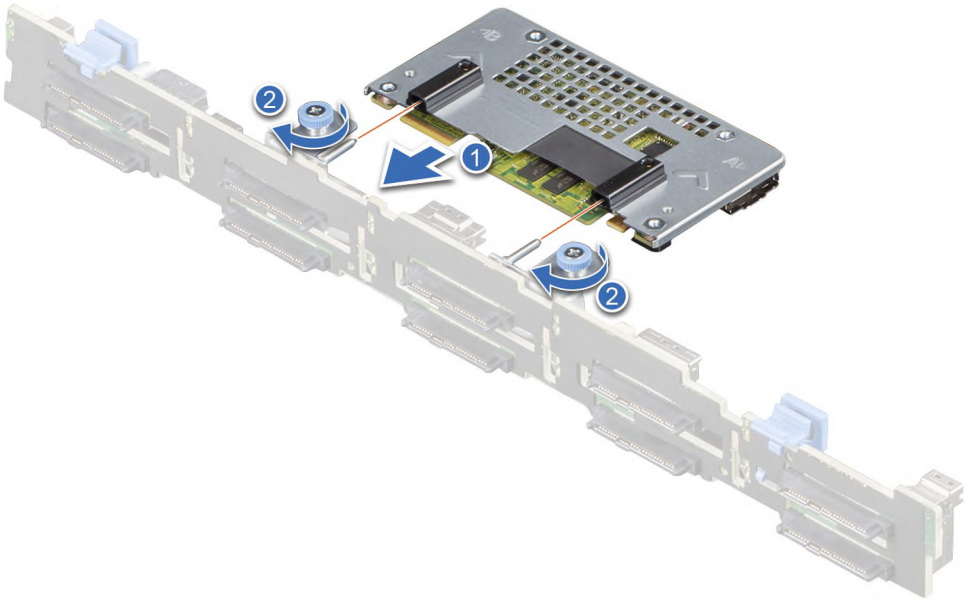

Figure 23. Install the PERC H755 front SAS card

Remove the PERC H350 adapter SAS

Prerequisites

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

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

Steps

1. 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 Install the PERC H350 adapter.
7. Close the system.
8. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.

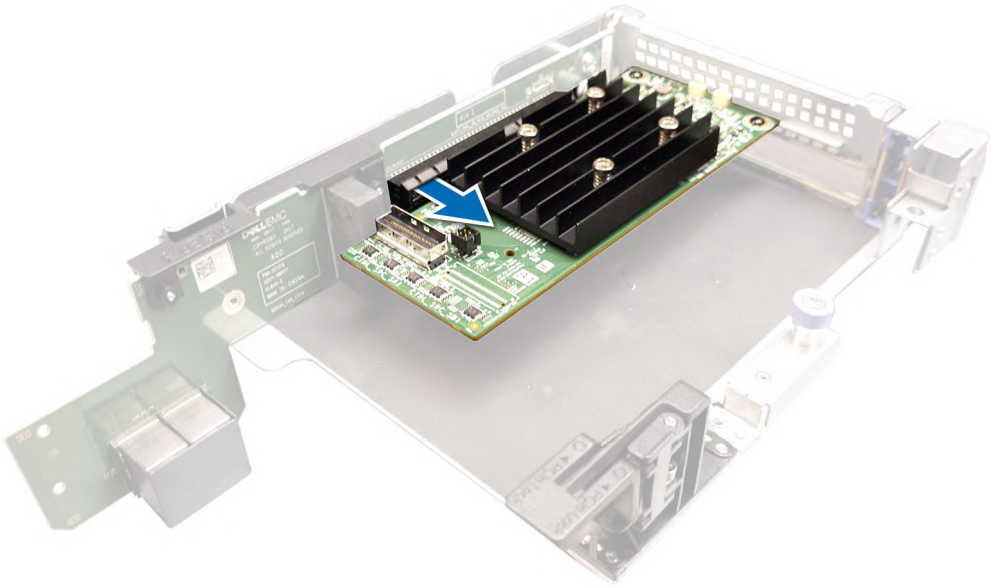

Figure 24. Remove the PERC H350 adapter SAS

Install the PERC H350 adapter SAS

Prerequisites

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

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

Steps

1. 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 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.

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

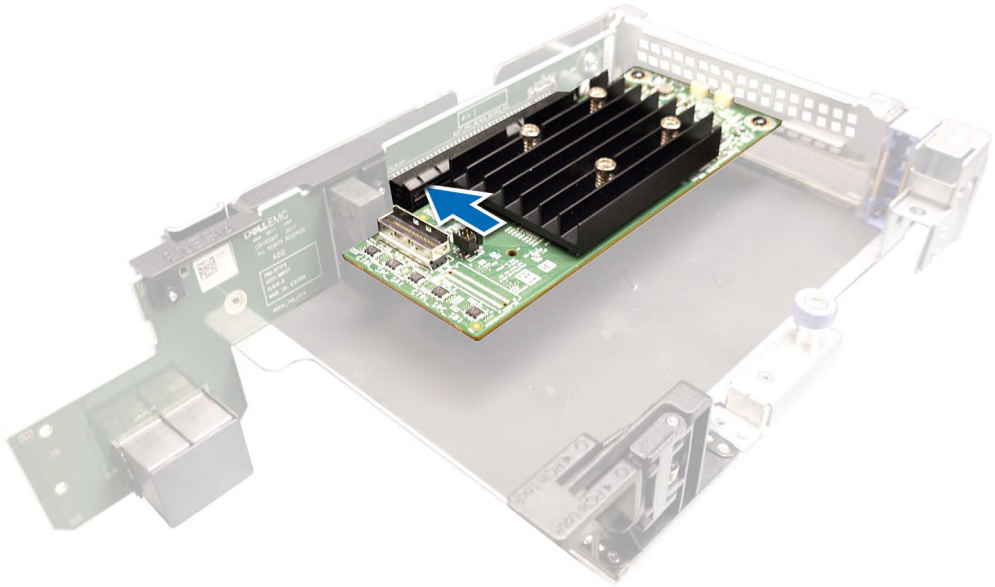

Figure 25. Install the PERC H350 adapter SAS

Remove PERC H350 Mini Monolithic SAS

Prerequisites

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

Steps

1. Using Phillips #2 screwdriver, loosen the screws that secure the storage controller cable to the connector on the system board.
2. Lift the storage controller cable to disconnect it from the connector on the system board.

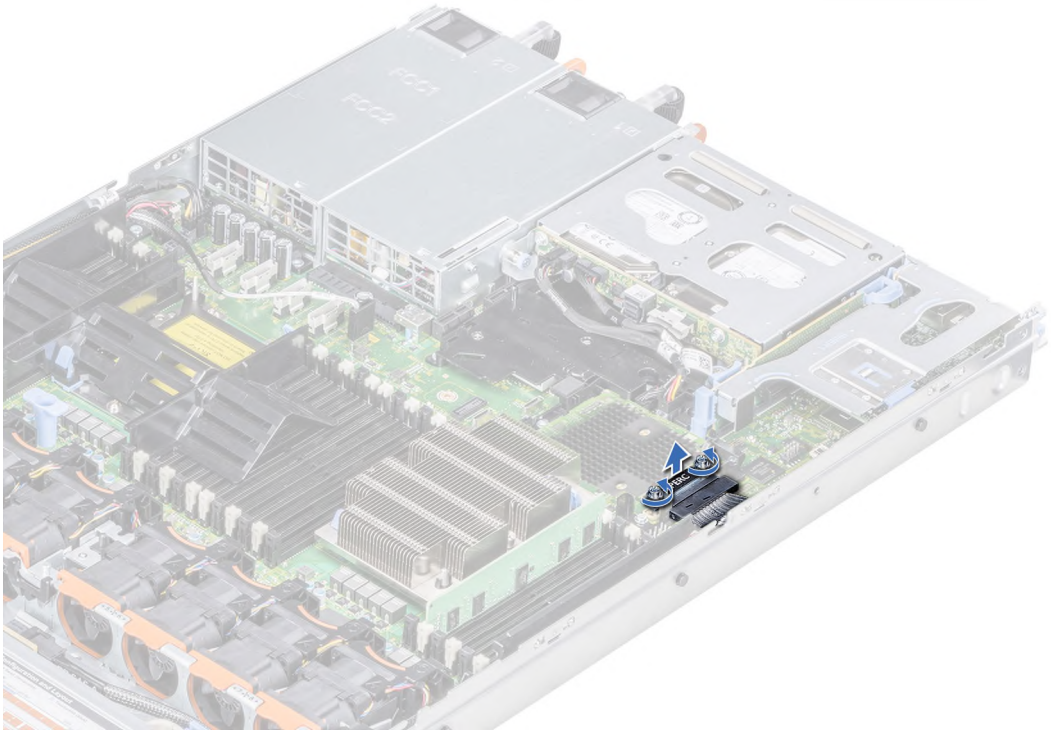

Figure 26. Remove the cable

3. Lift one end of the card and angle it to disengage the card from the card holder on the system board.
4. Lift the card out of the system.

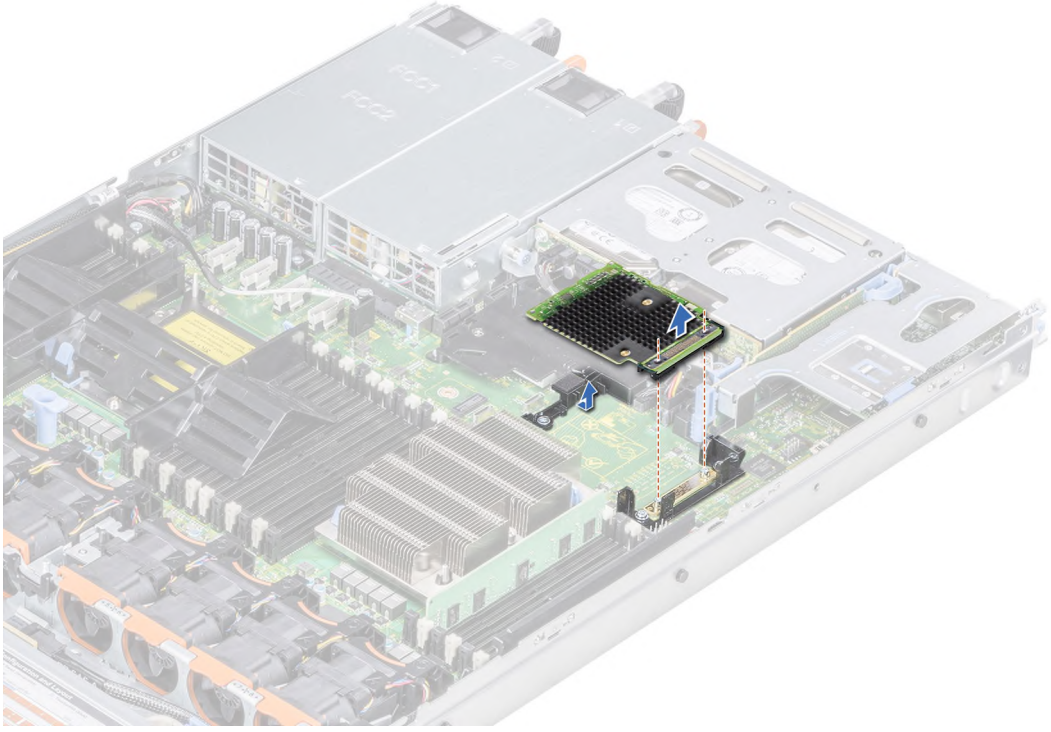

Figure 27. Remove the PERC H350 Mini Monolithic SAS

Install PERC H350 Mini Monolithic SAS

Prerequisites

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

Steps

1. Angle the integrated storage controller card and align the end of the card with the storage controller card connector on the system board.
2. Lower the connector side of the storage controller card into the storage controller card connector on the system board.

   NOTE: Ensure that the slots on the system board align with the screw holes on the storage controller card connector.

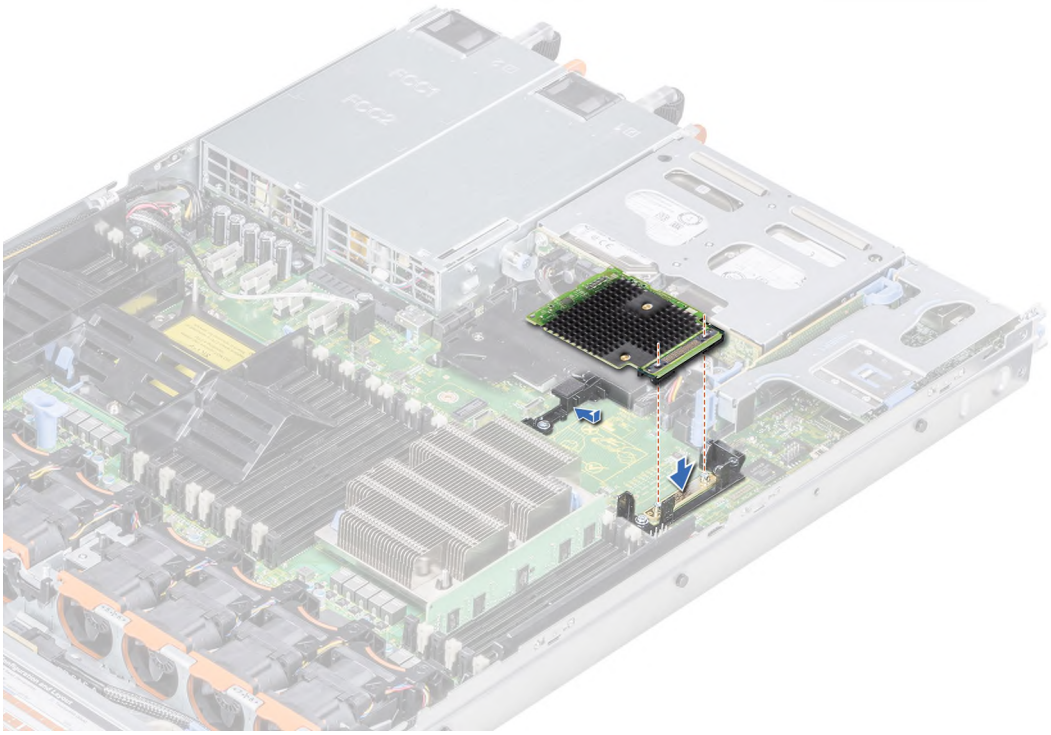

Figure 28. Install PERC H350 Mini Monolithic SAS

3. Route the storage controller card cable along with the wall of the system.
4. Align the screws on the integrated storage controller card cable with the screw holes on the connector.
5. Using Phillips #2 screwdriver, tighten the screws to secure the integrated storage controller card cable to the card connector on the system board.

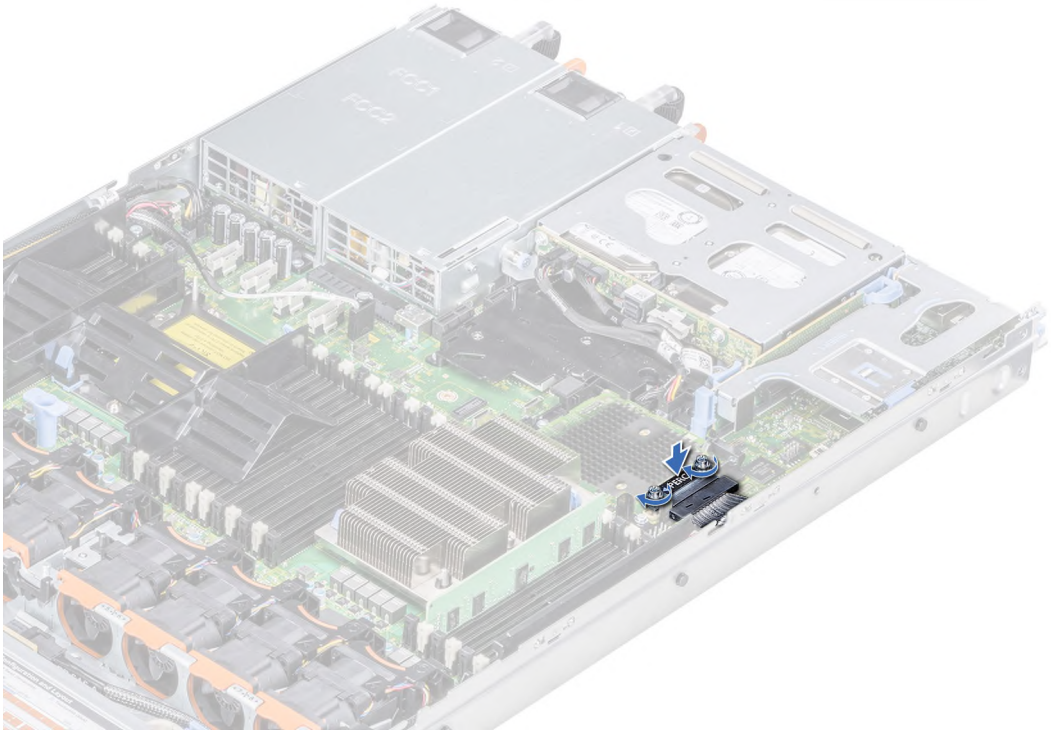

Figure 29. Install the cable

5

Driver support for PERC 11

The PERC 11 cards require software drivers to operate with the supported operating systems.

This chapter contains the procedures for installing the drivers for the PERC 11 cards.

NOTE: The driver for PERC 11 for VMware ESXi is packaged within the VMware ESXi ISO image that is downloaded from Dell. For more information, see the VMware documentation at Virtualization Solutions Documentation. It is not recommended to have drivers from controllers prior to PERC 11 on the same system.

The two methods for installing a driver that is discussed in this chapter are:

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

Topics:

- Creating the device driver media
- Windows driver installation
- Linux driver installation
- Load the driver while installing an operating system

Creating the device driver media

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

- Downloading Drivers From The Dell Support Website
- Downloading Drivers From The Dell Systems Service And Diagnostic Tools Media

Download and save PERC 11 drivers from the support site

About this task

To download drivers from the Dell Support website:

Steps

1. Go to the Support Site.
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 to load the driver. For more information about reinstalling the operating system, see the relevant section for your operating system later in this guide.

Download and save PERC 11 drivers from the Dell Systems Service and Diagnostic Tools

About this task

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

Steps

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

Windows driver installation

Before you install the Windows driver for PERC 11, 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 Support Site.
- Create a device driver media using one of the methods listed below:
  - USB drive
  - CD
  - DVD

Install PERC 11 driver while newly installing the Windows Server 2016 and later

About this task

To install the driver:

Steps

1. Boot the system using the Windows Server 2016, or newer media.
2. Follow the on-screen instructions until you reach **Where do you want to install Windows Server 2016 or later** window and then select **Load driver**.
3. As prompted, insert the installation media and browse to the appropriate location.
4. Select a PERC 11 series card from the list.
5. Click **Next** and continue installation.

Install PERC 11 driver on which the Windows Server 2016 is already installed and later

About this task

Perform the following steps to configure the driver for the RAID controller on which the Windows Server 2016 is already installed:

Steps

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 Install and remove a PERC 11 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.

Update PERC 11 driver that runs on Windows Server 2016 or later

Prerequisites

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

Steps

1. Insert the media containing the driver.
2. Select **Start** > **Settings** > **Control Panel** > **System.**
   The **System Properties** screen is displayed.
   NOTE: The path to **System** 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 **Storage Controllers** by double-clicking the entry or by clicking on the plus (+) symbol next to **Storage Controllers**.
6. Double-click the controller for which you want to update the driver.
7. Click the **Driver** tab and click **Update Driver**.
   The screen to update the device driver wizard is displayed.
8. Select **Install from a list or specific location**.
9. Click **Next**.
10. Follow the steps in the wizard and browse to the location of the driver files.
11. Select the INF file from the drive media.
12. Click **Next** and continue the installation steps in the wizard.
13. Click **Finish** to exit the wizard and reboot the system for the changes to take place.
    NOTE: Dell provides the Dell Update Package (DUP) to update drivers on systems running Windows Server 2016 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 Support Site.

Linux driver installation

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, see, Installing or updating the RPM driver package with KMOD support. If not, proceed with using the native device driver and then skip to the topic Installing or Updating the RPM Driver Package With KMP Support.

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.

Install or update a RPM driver package using the KMOD support

Prerequisites

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

About this task

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

Steps

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 the loaded driver version by running the following command: `modinfo megaraid_sas`.

Install or update a RPM driver package using the KMP support

Prerequisites

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

About this task

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

Steps

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 the loaded driver version by running the following command: `modinfo megaraid_sas`.

Upgrading the Kernel

About this task

When upgrading to a new kernel, you must reinstall the DKMS-enabled driver packages. Perform the following steps to update or install the driver for a new kernel:

Steps

1. At a **terminal** window, type the following: `dkms build -m <module_name> - v <module version> - k <kernel version> dkms install -m <module_name> - v <module version> - k <kernel version>`.
2. To check if the driver is successfully installed in the new kernel, type: `dkms status`.
   A message similar to the following is displayed: `<driver name>, <driver version>, <new kernel version>: installed`.
3. If the previous device driver is in use, you must restart the system for the updated driver to take effect.

Load the driver while installing an operating system

Steps

1. Perform the following operations to install driver media:
   - PERC Linux driver ISO:
     a. Download the PERC Linux driver package from the Dell Support site.
     b. Extract two base directories from the tar.gz package (**tar.gz** > **tar** > **base directories**).
     c. Extract the ISO file that is available in the zipped **disks-x** directory. For example, **RHEL79/disks-1/megaraid_sas-07.719.03.00_el7.9-1.x86_64.iso.gz > megaraid_sas-07.719.03.00_el7.9-1.x86_64.iso**
     d. Mount the ISO to the Server, burn the ISO to a CD or DVD or copy the ISO file to a USB. The USB has to match with the ISO.
   - LC driver pack:
     a. Install the LC driver pack.
     b. 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 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.

6

Firmware

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

Topics:

- Upgrade firmware controller using Dell Update Package (DUP)

Upgrade firmware controller using Dell Update Package (DUP)

About this task

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

Steps

1. Go to the Drivers and Downloads page on the support site.
2. Locate your controller.
3. Download the DUP file.
   a. To upgrade by using Windows or iDRAC, download the Windows executable file.
   b. To upgrade using Linux, download the **.bin** file.

   NOTE: For VMware, firmware must be upgraded by using iDRAC or the PERC CLI.
4. Install the DUP by doing one of the following:
   a. For Windows, run the executable file in the Windows environment.
   b. For Linux, run the **.bin** file in the Linux environment.
   c. For iDRAC, click **System iDRAC** > **Maintenance** > **System Update**, upload Windows executable, and then install.

7

Manage PERC 11 controllers using HII configuration utility

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

Topics:

- Enter the PERC 11 HII configuration utility
- Exit the PERC 11 HII configuration utility
- Navigate to Dell PERC 11 configuration utility
- View the HII Configuration utility dashboard
- Configuration management
- Controller management
- Virtual disk management
- Physical disk management
- Hardware components
- Security key management in HII configuration utility

Enter the PERC 11 HII configuration utility

About this task

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

Steps

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.

Exit the PERC 11 HII configuration utility

About this task

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

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.

Navigate to Dell PERC 11 configuration utility

Steps

1. Enter the UEFI configuration Utility. See Enter the PERC 11 HII configuration utility.
   The **Device Settings** screen displays a list of NIC ports and the RAID controllers.
2. To enter PERC 11 configuration utility, click the appropriate PERC controllers.
   The **Dashboard view** screen is displayed.

View the HII Configuration utility dashboard

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 9. Dashboard view screen

Dashboard view options	Description
Main menu	Displays the following configuration options: • 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: • 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: • Controller Management • Hardware Components • Physical Disk Management • Virtual Disk Management
Actions	Displays the following options: • 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

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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.

Create virtual disks

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See, Navigate to Dell PERC 11 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:

Table 10. Create virtual disks

Option	Description
Create Virtual Disk	Allows you to create virtual disk selecting the RAID level, physical disks, and 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.
   NOTE: Ensure that you restart the system after creating a new Non-RAID or Virtual Disk on drives that previously had boot partitions.

Configure virtual disk parameters

Steps

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 11. Configure virtual disk parameters

Virtual disk parameters	Description
Virtual Disk Name	Allows you to enter the name for the virtual disk **NOTE:** Allowed characters are A-Z, a-z, 0-9, underscore (_), and hyphen (-) only.
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: ● 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: ● 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.
Default Initialization	Displays the virtual disk initialization options. You can set the default initialization to: ● 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. For more information, see Virtual disk initialization. By default, the default initialization is set to No.

Create profile based virtual disk

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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.

View disk group properties

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
2. Click **Main Menu** > **Configuration Management** > **View Disk Group Properties**.
   The list of disk group properties are displayed:

Table 12. View disk group properties

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

Convert to Non–RAID disk

Prerequisites

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

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
2. Click **Main Menu** > **Configuration Management** > **Convert to Non–RAID Disk**.
   The list of physical disks appears.
3. Select the physical disk to convert to Non–RAID disk.
4. Click **Ok**.
   A screen appears asking if you are sure you want to perform the operation.
5. Select the **Confirm** option.
6. Click **Yes**.
   The operation is successful.

Delete configurations

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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. **CAUTION: It is recommended that you back up data stored on the virtual disks and hot spare disks on the controller before deleting the virtual drive.**

   Select **Confirm** and click **Yes** to continue.
   The virtual disks and hot spare disks available on the controller are deleted successfully.

Controller management

Clear controller events

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
2. Click **Main Menu** > **Controller Management** > **Advanced Controller Management**.
3. Click **Clear Controller Events**.
   A screen is displayed asking if you are sure you want to clear the controller events.
4. Select **Confirm** and click **Yes** to continue.

Save controller events

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
2. Click **Main Menu** > **Controller Management** > **Advanced Controller Management**.
3. Click **Save Controller Events**.
   A screen is displayed asking if you want to replace the existing file name.
4. Select **Confirm** and click **Yes** to continue.

Save debug log

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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 successful.
4. Click **Ok**.

Enable security

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
2. Click **Main Menu** > **Controller Management** > **Advanced Controller Management**.
3. Click **Enable security**, select **Local Key Management**.
4. Click **Ok**.
5. If you want to use the passphrase generated by the controller, click **Suggest Passphrase** and **Confirm** the passphrase by re-entering.
   The operation is successful.
6. Select **I Recorded the Security Settings For Future Reference**, click **Enable Security**.
   A screen is displayed indicating that the security will be enabled on this controller if you proceed.
7. Select **Confirm** and click **Yes** to continue.
   The operation is successful and click **Ok**.

Disable security

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.

2. Click **Main Menu** > **Controller Management** > **Advanced Controller Management**.
3. Click **Disable security**.
   A screen is displayed asking if you are sure you want to disable security.
4. Select **Confirm** and click **Yes** to continue.
   The operation is successful and click **Ok**.

Change security settings

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
2. Click **Main Menu** > **Controller Management** > **Advanced Controller Management**.
3. Click **Change Security Settings**, select **Change Current Security Settings**.
4. Click **Ok**.
5. If you want to use the passphrase generated by the controller, click **Suggest Passphrase** and **Confirm** the passphrase by re-entering.
   The operation is successful.
6. Click **Save Security Settings**.
7. Select **Confirm** and click **Yes** to continue.
   The operation is successful and click **Ok**.

Restore factory default settings

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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.

Auto configure behavior

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
2. Click **Main Menu** > **Controller Management** > **Advanced Controller Management > Manage Controller Mode**.
   You can view the current Controller Mode.
3. Click **Manage Controller Mode**.
   If required, you can view or change the hard drive settings for the controller. The possible options are:
   - **Off** and **Non–RAID Disk**
4. Click **Apply Changes** to save the changes.
5. Select **Confirm** and click **Yes** to continue.

   NOTE: This feature is supported on PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 Mini Monolithic SAS, and PERC H350 adapter SAS.

Manage controller profile

About this task

View the details of the profile and choose the desired profile, if supported. To view the properties of the controller profiles:

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.

2. Click **Main Menu** > **Controller Management** > **Advanced Controller Management** > **Manage Controller Profiles**.
   The current profile and profile properties are displayed.

Advanced controller properties

Set the patrol read mode

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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**.

Enable physical disk power management

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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.

Configure hot spare drives

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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.

Set task rates

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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.

Properties of Enterprise Key Management (EKM)

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
2. Click **Main Menu** > **Controller Management** > **Advanced Controller Properties**.
3. Click **Enterprise Key Management**.
   The properties of Enterprise Key Management is displayed.

Controller properties

Auto import foreign configuration

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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** or **Disabled**.
4. Click **Apply Changes**.

Disable auto import

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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.

Enable auto import

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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.

Select boot mode

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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 13. 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.

   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.

Abort the consistency check

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
2. Click **Main Menu** > **Controller Management** > **Advanced Controller Properties**.
3. In the **Controller Properties** section, set the **Abort Consistency Check on Error** option to **Enabled** or **Disabled**.
4. Click **Apply Changes**.
   The option to abort the consistency check operation on a redundant virtual disk is enabled if there is any inconsistency found in the data.

Preboot trace buffer

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
2. Click **Main Menu** > **Controller Management** > **Advanced Controller Properties**.
3. In the **Controller Properties** section, set the **Preboot Trace Buffer** option to **Enabled** or **Disabled**.
4. Click **Apply Changes**.

Clear the cache memory

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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.

Enable boot support

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
2. Click **Main Menu** > **Controller Management**.
3. From the **Select Boot Device** drop-down box, select the primary bootable device.

   In **Select Boot Device**, you will not be able to view 4 K sector drives. To view all the virtual disks created, navigate to the **Virtual Disk Management** screen in HII. For more information, see Virtual disk management.

   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 is auto-selected as the boot device, if the controller does not have any virtual disks present.

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

   NOTE: 4 K sector drives boot support is only available in UEFI mode and managed by the boot loader.
4. Click **Apply Changes**.
   Boot support is enabled for the selected controller.

Virtual disk management

Virtual disk numbering

Virtual disks are numbered in descending order beginning with the highest, ID 239.

View virtual disk properties

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
2. Click **Main Menu** > **Virtual Disk Management**.
   All the virtual disks associated with the RAID controller are displayed.
3. To view the properties, click on the virtual disk. You can view the following properties of the Virtual disk:

Table 14. Virtual disk properties

Option	Description
Operation	List of operations you can perform on the selected virtual disk. The options are: • Blink • Unblink • Delete Virtual Disk • Reconfigure Virtual Disks • Fast Initialization • Slow Initialization
Name	Indicates the name of the virtual disk.
RAID level	Indicates the RAID level of the virtual disk.
Status	Indicates the status of the virtual disk. The possible options are: • Optimal • Degraded • Offline • Failed
Size	Indicates the size of the virtual disk.

View physical disks associated with a virtual disk

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
2. Click **Main Menu** > **Virtual Disk Management**.
   All the virtual disks associated with the RAID controller are displayed.
3. Click on a virtual disk.
   The properties of the virtual disk are displayed.
4. Click **View Associated Physical Disks**.
   All the physical disks that are associated with the virtual disk are displayed.
5. From the **Associated Physical Disks** section, select the physical disk.
6. Click **View Physical Disk Properties** to view the physical disk properties.

View advanced properties of a virtual disk

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
2. Click **Main Menu** > **Virtual Disk Management**.
   All the virtual disks associated with the RAID controller are displayed.
3. Click the virtual disk.
   The properties of the virtual disk are displayed.
4. Click **Advanced...**.
   You can view the following additional properties of the virtual disk:

Table 15. 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.
Secured	Indicates whether the virtual disk is secured or not.
Bad blocks	Indicates whether the virtual disk has corrupted blocks.

Configure virtual disk policies

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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 16. 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: • 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:

Table 16. Virtual disk policies (continued)

Option	Description
	• No Read Ahead • Read Ahead
Disk cache	Allows selection of the disk cache policy for the virtual disk. The possible options are: • Default (Disk Default) • Enable • Disable

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

Configure Virtual Disks

When configuring the virtual disks, you should consider the workload intended; RAID 1: for simple boot disk; RAID 5 or 6: for file or web servers (sequential reads/writes of files); RAID 10: 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.

Perform expand virtual disk operation

Prerequisites

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

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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.

Perform consistency check

Prerequisites

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

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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

View physical disk properties

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
2. Click **Main Menu** > **Physical Disk Management**.
   All the physical disks that are associated with the RAID controller are displayed.
3. To view the properties, click the physical disk.

Table 17. Physical disk properties

Option	Description
Operation	The list of operations you can perform on the selected physical disk. The options are: • Blink • Unblink • Assign global hot spare • Cryptographic erase • Convert to non–RAID disk
Device ID	Unique identifier of the physical disk.
Backplane ID	Backplane ID in which the physical disk is located in for PERC H755 adapter, PERC H755 front SAS, PERC H755N front NVMe, PERC H750 adapter SAS, PERC H755 MX adapter, PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS , and PERC H350 Mini Monolithic SAS
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.

Table 17. Physical disk properties (continued)

Option	Description
Serial number	Serial of the physical disk.

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

Table 18. 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
SMART status	SMART status of a physical disk
Revision	Firmware version of the physical disk
WWID	Unique identifier used to identify the device
Multipath	Multipath of the controller
Physical disk power state	Power condition (On or Power Save) of the physical disk
Disk cache setting	Disk cache setting **NOTE:** Disk cache for SATA Gen3 drives is disabled by default.
Disk protocol	Type of hard disk used
Device speed	Speed of the physical disk
Negotiated link speed	Negotiated link speed of the device
PCIe capable link width	N/A for SAS/SATA drives
PCIe negotiated link width	N/A for SAS/SATA drives
Encryption capable	Encryption capability of the physical disk
Encryption supported	Encryption capability enabled at the controller level
Secured	Security status of the physical disk
Cryptographic erase capable	Cryptographic erase capability of the physical disk

Cryptographic erase

Cryptographic erase is a process to erase all data permanently on an encryption-capable and unconfigured physical disk, and reset the security attributes.

Prerequisites

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

About this task

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

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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

Prerequisites

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

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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 a global hot spare

Prerequisites

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

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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 a dedicated hot spare

Prerequisites

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

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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

Prerequisites

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

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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 operation is successful.

Convert to Non–RAID disk

Prerequisites

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

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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 operation is successful.

Hardware components

View battery properties

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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 19. 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 20. 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.

View physical disks associated with an enclosure

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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 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 Applications and User Interfaces supported by PERC 11 .

- 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 or changes the local key management (LKM) security key. Deletes the local key management (LKM) or secure enterprise key manager (SEKM) security key.
- The **Virtual Disk Management** screen displays physical disk information and action menus. You can perform the following security related actions through the virtual disk management menu:
  - **Secure Disk Group**—Secures all virtual disks in disk group.
  - **Create secure virtual disk**—Creates a new virtual disk that is secured with the security key on the controller.
- The **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:
  - **Secure non–RAID disk**—Secures the non–RAID disk with the controller security key.
  - **Cryptographic Erase**—Permanently erases all data on the physical disk and resets the security attributes.

For more information on the Physical Disk Management screen and the Virtual Disk Management screen, see Physical disk management and Virtual disk management.

8

Security key and RAID management

Topics:

- Security key implementation
- Local Key Management
- Create a security key
- Change Security Settings
- Disable security key
- Create a secured virtual disk
- Secure a non-RAID disk
- Secure a pre-existing virtual disk
- Import a secured non-RAID disk
- Import a secured virtual disk
- Dell Technologies OpenManage Secure Enterprise Key Manager

Security key implementation

The PERC 11 series of cards support self-encrypting disk (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 using local key management (LKM) or OpenManage Secure Enterprise Key Manager, also referred as Secure Enterprise Key Manager (SEKM). The LKM 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.

NOTE: If the host system is powered off when connected to an external enclosures or if the sled is powered off in C6XXX PowerEdge servers, the drives will remain in an unlocked state until they are power cycled or AC power is disconnected from the sled or external enclosure.

Local Key Management

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

NOTE: LKM mode is not supported on PERC H355 adapter SAS, PERC H350 adapter SAS, PERC H355 front SAS, and PERC H350 Mini Monolithic SAS.

Create a security key

About this task

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

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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.

Change Security Settings

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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.

Disable security key

About this task

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

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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 a secured virtual disk

About this task

To create a secured virtual disk, the controller must have a security key established first. See Create a security key.

NOTE: Combining SAS and SATA hard drives within a virtual disk is not supported. Also, combining hard drives and olid-state drives (SSDs) within a virtual disk is not supported. Mixing of NVMe drives is not supported.

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

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
2. Click **Main Menu** > **Configuration Management** > **Create Virtual Disk**.
   For more information, see Create virtual disks.
3. Select the **Secure Virtual Disk** option.
4. Click **Create Virtual Disk**.
   The secure virtual disk is created successfully.

Secure a non-RAID disk

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

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
2. Click **Main Menu** > **Physical Disk Management**.
   The list of Non-RAID disks is displayed.
3. Select a non-RAID disk.
4. From the **Operations** drop-down menu, select **Secure Non-RAID Disk**.

Secure a pre-existing virtual disk

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
2. Click **Main Menu** > **Virtual Disk Management**.
   The list of virtual disks is displayed.
3. Select a virtual disk.
4. From the **Operations** drop-down menu, select **Secure Virtual Disk**.

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

Import a secured non-RAID disk

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

Prerequisites

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

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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 non-RAID disk with a different passphrase.
5. Select the **Confirm** option.
6. Click **Yes**.

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

Import a secured virtual disk

Prerequisites

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

Steps

1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 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.

Dell Technologies 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 H755 adapter, PERC H755 front SAS, and PERC H755N front NVMe, and allows the creation of secured virtual disks and non–RAID disks. For more information about supported platforms, see Support Site.

Enterprise key manager mode is not supported on the PERC H755 MX adapter, PERC H355 front SAS, PERC H355 adapter SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS.

Manage enterprise key manager mode

iDRAC manages Enterprise key manager features. For instructions on enabling enterprise key manager mode, vidi **dell.com/idracmanuals**.

NOTE: If preserved cache is present, the controller does not allow OpenManage Secure Enterprise Key Manager (SEKM) mode to be enabled.

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: Transitioning a controller from Local Key Management (LKM) mode to SEKM mode is supported on firmware starting with version 52.16.1-4074.

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

Disable enterprise key manager mode

Enterprise key manager mode can be disabled from any supported Applications & User Interfaces supported by PERC 11. For more information, see the management application's user's guide or see Disable security key.

Manage 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 Create a secured virtual disk.

Manage non–RAID disks in enterprise key manager mode

Non–RAID disks are managed in the same way in enterprise key manager mode as in local key manager mode. SED capable non–RAID disks can be secured after creation. See Create a secured virtual disk.

Transition of drives from local key management to enterprise key management (without supported firmware for PERC and iDRAC)

Local key management drives can be transitioned 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 transition from local key management drives to enterprise key management:

Steps

1. Save the current local key management security key.
2. Shut down both systems.
3. Remove the local key management drives and reinsert them to the enterprise key manager enabled system.
4. Power on the enterprise key manager system.
5. Go 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 about performing this action, contact Support Site.

Migrate of drives from local key management to enterprise key management (with supported firmware for PERC and iDRAC)

PERC enables transition from Local Key Management (LKM) mode to Secure Enterprise Key Manager (SEKM) mode without disabling LKM security first. For instructions on transitioning from LKM mode to SEKM mode, see iDRAC Manuals.

NOTE: This feature is supported on firmware starting with version 51.16.0-4076.

The transition from LKM to SEKM on the controller fails if the following are true at time of attempt:

- Snapdump is present on PERC.
- Preserved cache is present on PERC.
- RAID level migration is in progress on PERC,
- Online capacity expansion is in progress on PERC.
- Sanitize on a physical disk is in progress.
- LKM key that does not match with the current key of PERC.
- PERC firmware does not support transition.

9

Troubleshooting issues in PERC11 cards

To get help for resolving issues in your PERC11 series cards, you can contact your Dell Technical Service representative.

Topics:

- Single virtual disk performance or latency in hypervisor configurations
- 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
- Foreign configuration found error message
- Foreign configuration not found in HII
- 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
- World Wide Number on drive sticker is not the same in applications
- Backplane firmware revision not changing in PERC interfaces after an update

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 under lying 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.

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 cables from the PERC controller to the backplane 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 cables from the PERC controller to the backplane might be improperly connected.
Corrective Action:	Check the cable connections and fix any problems. Restart the system. Use the HII configuration utility to import the virtual disk or discard the preserved cache. For the steps to discard the preserved cache, see Clear 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 cables from the PERC controller to the backplane might be improperly connected.
Corrective Action:	Check the cable connections and fix any problems. Restart the system.

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 cables from the PERC controller to the backplane 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.

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:	`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 Clear the cache memory.

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

Security key errors

Secured foreign import errors

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

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

- **The 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 Create 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 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.

You can perform Cryptographic Erase only on encryption-capable disks that are not hot spares and not configured as non-RAID or virtual disks. Ensure that the conditions are met and see Cryptographic Erase.

General issues

PERC card has yellow bang in Windows operating system 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 support for PERC 11.

PERC card not seen in operating systems

Issue:	The device does not appear in the **Device Manager**.
Corrective Action:	Power off the system and reseat the controller.

Issues in controller, battery, and disk when operating at low temperature

Issue:	If the controller is operating at temperatures less than zero degree Centigrade, then an increase in the number of issues related to controller, battery, or drive is observed.
Corrective Action:	Ensure that the controller ambient temperature is more than zero degree Centigrade.

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 the support site 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 Global Technical Support.

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 reseat 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 I/O stress. There is only one rebuild I/O operation for every five host I/O operations.
Corrective Action:	If possible, reduce I/O stress on the physical disk or increase the value of rebuild rate controller parameter.

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

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

Description:	This transient state could be reported as an event in management applications as **A foreign configuration was detected on RAID Controller is SL x**, where x is the slot of the RAID controller.
Corrective Action:	No action is required on the foreign configuration state of the drive as it is transient and the controller handles the event automatically.

SMART errors

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

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

Smart error detected on a non–RAID disk

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

Smart error detected on a physical disk in a non–redundant virtual disk

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

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

Issue:	A SMART error is detected on a physical disk in a redundant virtual disk.
Corrective Action:	Perform the following steps: 1. Back up your data. 2. Force the physical disk offline. **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 Configure hot spare drives.

Replace member errors

NOTE: For more information about the **Replace Member** features, see Configure hot spare drives.

Source disk fails during replace member operation

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

Target disk fails during replace member operation

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

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

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

Linux operating system errors

Virtual disk policy is assumed as write-through

Error:	`<Date:Time> <HostName> kernel: sdb: asking for cache data failed<Date:Time> <HostName> kernel: sdb: assuming drive cache: write through`
Corrective Action:	The error message is displayed when the Linux Small Computer System Interface (SCSI) mid-layer asks for physical disk cache settings. The controller firmware manages the virtual disk cache settings on a per controller and a per virtual disk basis, so the firmware does not respond to this command. The Linux SCSI mid-layer assumes that the virtual disk's cache policy is **Write-Through**. SDB is the device node for a virtual disk. This value changes for each virtual disk. 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 30. 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 21. Drive indicator codes

Drive status indicator code	Condition
Blinks green twice per second	The drive is being identified or preparing for removal
Off	The drive is ready for removal **NOTE:** The drive status indicator remains off until all drives are initialized after the system is powered on. Drives are not ready for removal during this time.
Blinks green, amber, and then powers off	There is an expected drive failure
Blinks amber four times per second	The drive has failed
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 **BP_PSV** and **Firmware version** is displayed as **03**.
Corrective action	Turn off the system, reseat 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).

World Wide Number on drive sticker is not the same in applications

World Wide Number (WWN) on the drive sticker and applications are not matching.	NVMe drives do not have a WWN. So, the applications create a WWN from the available drive information. This WWN may not match with the WWN that is on the drive sticker, if present.

Backplane firmware revision not changing in PERC interfaces after an update

After updating the backplane firmware on 15G and later PowerEdge servers, the backplane version will not show as updated on some interfaces until the system is reset.

10

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 11 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 22. 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 23. 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 23. 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.

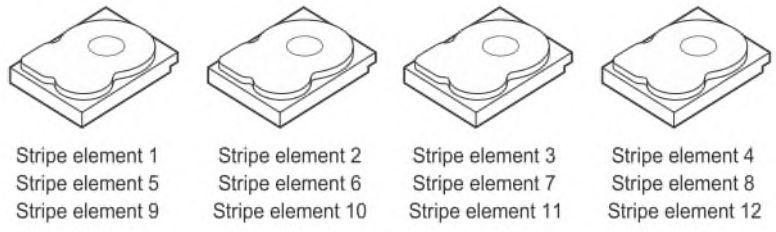

Figure 31. Example of disk striping (RAID 0)

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.

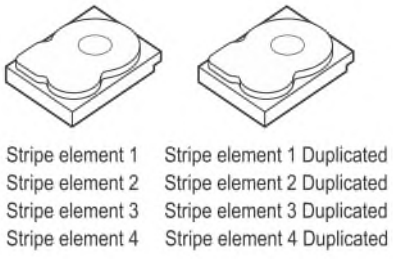

Figure 32. Example of Disk Mirroring (RAID 1)

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.

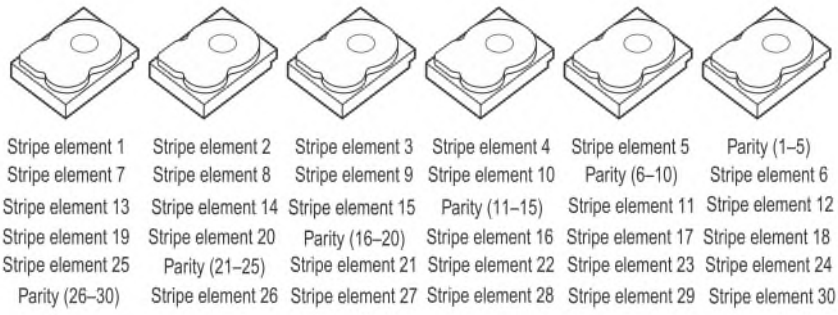

Figure 33. Example of Distributed Parity (RAID 5)

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

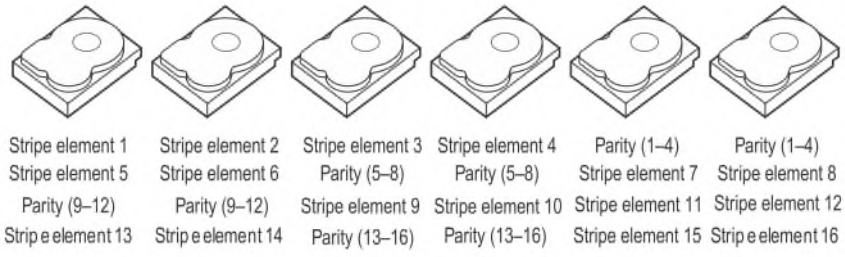

Figure 34. Example of Dual Distributed Parity (RAID 6)

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

11

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 the How to Recycle page and select the relevant country.

Contacting Dell

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

Steps

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

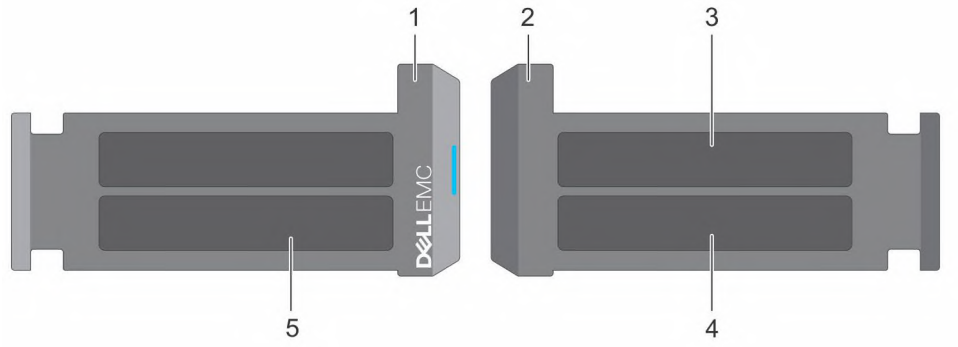

Figure 35. Locating the Express Service Code and Service tag

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

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

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

Receiving automated support with SupportAssist

Dell SupportAssist is an optional Dell Services offering that automates technical support for your Dell 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 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 Technical Support.
- Automated diagnostic collection — SupportAssist automatically collects system state information from your devices and uploads it securely to Dell. This information is used by Dell Technical Support to troubleshoot the issue.
- Proactive contact — A Dell Technical Support agent contacts you about the support case and helps you resolve the issue.

The available benefits vary depending on the Dell Service entitlement purchased for your device. For more information about SupportAssist, go to the SupportAssist page.

12

Documentation resources

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

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

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

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

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

Table 24. Additional documentation resources for your system

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

Table 24. Additional documentation resources for your system (continued)

Task	Document	Location
	? > **About**.	
	For information about updating drivers and firmware, see the Methods to download firmware and drivers section in this document.	Drivers
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**.	PowerEdge Server Event and Error Messages