

3. To update the plugin, click **Update Plugin**.
In the **Confirmation** window, select the **I agree that I have captured a snapshot of the OpenManage Enterprise appliance prior to performing a plugin action** option, and then click **Update**.

Results

After update operation is complete, the version is displayed in the plugin section.

Execute remote commands and scripts

About this task

When you get an SNMP trap, you can run a script on OpenManage Enterprise. This sets up a policy that opens a ticket on your third party ticketing system for alert management. You can create and store only up to **four** remote commands.

① **NOTE:** The use of the following special characters as RACADM and IPMI CLI parameters is not supported: [, ; , | , \$, > , < , & , ' ,] , . , * , and ' .

Steps

1. Click **Application Settings > Script Execution**.

2. In the **Remote Command Setting** section, do the following:

- a. To add a remote command, click **Create**.
 - b. In the **Command Name** box, enter the command name.
 - c. Select any one of the following command type:
 - i. Script
 - ii. RACADM
 - iii. IPMI Tool
 - d. If you select **Script**, do the following:
 - i. In the **IP Address** box, enter the IP address.
 - ii. Select the authentication method: **Password** or **SSH Key**.
 - iii. Enter the **user name** and **password** or the **SSH Key**.
 - iv. In the **Command** box, type the commands.
 - Up to 100 commands can be typed with each command required to be on a new line.
 - Token substitution in scripts is possible. See *Token substitution in remote scripts and alert policy* on page 194
 - v. Click **Finish**.
 - e. If you select **RACADM**, do the following:
 - i. In the **Command Name** box, enter the command name.
 - ii. In the **Command** box, type the commands. Up to 100 commands can be typed with each command required to be on a new line.
 - iii. Click **Finish**.
 - f. If you select **IPMI Tool**, do the following:
 - i. In the **Command Name** box, enter the command name.
 - ii. In the **Command** box, type the commands. Up to 100 commands can be typed with each command required to be on a new line.
 - iii. Click **Finish**.
3. To edit a remote command setting, select the command, and then click **Edit**.
 4. To delete a remote command setting, select the command, and then click **Delete**.

OpenManage Mobile settings

OpenManage Mobile (OMM) is a systems management application that allows you to securely perform a subset of data center monitoring and remediation tasks on one or more OpenManage Enterprise consoles and/or integrated Dell Remote Access Controllers (iDRACs) by using your Android or iOS device. Using OMM you can:

- Receive alert notifications from OpenManage Enterprise.
- View the group, device, alert, and log information.

- Turn on, turn off, or restart a server.

By default, the push notifications are enabled for all alerts and critical alerts. This chapter provides information about the OMM settings that you can configure by using OpenManage Enterprise. It also provides information required to troubleshoot OMM.

NOTE: For information about installing and using OMM, see the *OpenManage Mobile User's Guide* at Dell.com/OpenManageManuals.

Related tasks

- Enable or disable alert notifications for OpenManage Mobile on page 186
- Enable or disable OpenManage Mobile subscribers on page 186
- Delete an OpenManage Mobile subscriber on page 187
- View the alert notification service status on page 187
- Troubleshooting OpenManage Mobile on page 189

Related information

- Enable or disable alert notifications for OpenManage Mobile on page 186
- Enable or disable OpenManage Mobile subscribers on page 186
- Troubleshooting OpenManage Mobile on page 189

Enable or disable alert notifications for OpenManage Mobile

About this task

By default, OpenManage Enterprise is configured to send alert notifications to the OpenManage Mobile application. However, alert notifications are sent from OpenManage Enterprise only when a OpenManage Mobile user adds OpenManage Enterprise to the OpenManage Mobile application.

NOTE: The administrator rights are required for enabling or disabling alert notifications for OpenManage Mobile.

NOTE: For OpenManage Enterprise to send alert notifications to OpenManage Mobile, ensure that the OpenManage Enterprise server has outbound (HTTPS) Internet access.

To enable or disable alert notifications from OpenManage Enterprise to OpenManage Mobile:

Steps

1. Click **OpenManage Enterprise > Application Settings > Mobile**.
2. Select the **Enable push notifications** check box.
3. Click **Apply**.

Related tasks

OpenManage Mobile settings on page 185

Related information

- OpenManage Mobile settings on page 185
- Delete an OpenManage Mobile subscriber on page 187

Enable or disable OpenManage Mobile subscribers

About this task

The check boxes in the **Enabled** column in the **Mobile Subscribers** list allow you to enable or disable transmission of alert notifications to the OpenManage Mobile subscribers.

NOTE:

- The administrator rights are required for enabling or disabling OpenManage Mobile subscribers.
- OpenManage Mobile subscribers may be automatically disabled by OpenManage Enterprise if their mobile service provider push notification service indicates that the device is permanently unreachable.

- Even if an OpenManage Mobile subscriber is enabled in the **Mobile Subscribers** list, they can disable receiving alert notifications in their OpenManage Mobile application settings.

To enable or disable alert notifications to the OpenManage Mobile subscribers:

Steps

1. Click **OpenManage Enterprise > Application Settings > Mobile**.
2. To enable, select the corresponding check box and click **Enable**. To disable, select the check box and click **Disable**.
You can select more than one subscriber at a time.

Related tasks

OpenManage Mobile settings on page 185

Related information


OpenManage Mobile settings on page 185

Delete an OpenManage Mobile subscriber on page 187

Delete an OpenManage Mobile subscriber

About this task

Deleting an OpenManage Mobile subscriber removes the user from the subscribers list, preventing the user from receiving alert notifications from OpenManage Enterprise. However, the OpenManage Mobile user can re-subscribe to alert notifications from the OpenManage Mobile application at a later time.

 **NOTE:** The administrator rights are required for deleting an OpenManage Mobile subscriber.

To delete an OpenManage Mobile subscriber:

Steps

1. Click **OpenManage Enterprise > Application Settings > Mobile**.
2. Select the check box corresponding to the subscriber name and click **Delete**.
3. When prompted, click **Yes**.

Related tasks

Enable or disable alert notifications for OpenManage Mobile on page 186

Enable or disable OpenManage Mobile subscribers on page 186

Delete an OpenManage Mobile subscriber on page 187

View the alert notification service status on page 187

Related information

OpenManage Mobile settings on page 185

Delete an OpenManage Mobile subscriber on page 187

View the alert notification service status

About this task

OpenManage Enterprise forwards alert notifications to OpenManage Mobile subscribers through their respective device platform alert notification service. If the OpenManage Mobile subscriber has failed to receive alert notifications, you can check the **Notification Service Status** to troubleshoot alert notification delivery.

To view the status of the alert notification service, click **Application Settings > Mobile**.

Related tasks

View the alert notification service status on page 187

Related information

OpenManage Mobile settings on page 185




Delete an OpenManage Mobile subscriber on page 187

View the alert notification service status on page 187

Notification service status

The following table provides information about the **Notification Service Status** displayed on the **Application Settings > Mobile** page.

Table 28. Notification service status

Status Icon	Status Description
	The service is running and operating normally. NOTE: This service status only reflects successful communication with the platform notification service. If the device of the subscriber is not connected to the Internet or a cellular data service, notifications will not be delivered until the connection is restored.
	The service experienced an error delivering a message which may be of a temporary nature. If the issue persists, follow troubleshooting procedures or contact technical support.
	The service experienced an error delivering a message. Follow troubleshooting procedures or contact technical support as necessary.

View information about OpenManage Mobile subscribers

About this task

After an OpenManage Mobile user successfully adds OpenManage Enterprise, the user is added to the **Mobile Subscribers** table in OpenManage Enterprise. To view information about the mobile subscribers, in OpenManage Enterprise, click **Application Settings > Mobile**.

You can also export the information about mobile subscribers to a .CSV file by using the **Export** drop-down list.

OpenManage Mobile subscriber information

The following table provides information about the **Mobile Subscribers** table displayed on the **Application Settings > Mobile** page.

Table 29. OpenManage Mobile subscriber information

Field	Description
ENABLED	Select or clear the check box, and then click Enable or Disable respectively to enable or disable the alert notifications to an OpenManage Mobile subscriber.
STATUS	Displays the status of the subscriber, indicating whether or not OpenManage Enterprise is able to send alert notifications successfully to the Alert Forwarding Service.
STATUS MESSAGE	Status description of the status message.

Table 29. OpenManage Mobile subscriber information (continued)

Field	Description
USER NAME	Name of the OpenManage Mobile user.
DEVICE ID	Unique identifier of the mobile device.
DESCRIPTION	Description about the mobile device.
FILTER	Filters are policies that the subscriber has configured for alert notifications.
LAST ERROR	The date and time the last error occurred when sending an alert notification to the OpenManage Mobile user.
LAST PUSH	The date and time the last alert notification was sent successfully from OpenManage Enterprise to the Alert Forwarding Service.
LAST CONNECTION	The date and time the user last accessed OpenManage Enterprise through OpenManage Mobile.
REGISTRATION	The date and time the user added OpenManage Enterprise in OpenManage Mobile.

Troubleshooting OpenManage Mobile

If OpenManage Enterprise is unable to register with the Message Forwarding Service or successfully forward notifications, the following resolutions are available:

Table 30. Troubleshooting OpenManage Mobile

Problem	Reason	Resolution
OpenManage Enterprise is unable to connect to the Dell Message Forwarding Service. [Code 1001/1002]	Outbound Internet (HTTPS) connectivity is lost.	By using a web browser, check if outbound Internet connectivity is available. If connection is unavailable, complete the following network troubleshooting tasks: <ul style="list-style-type: none"> • Verify if the network cables are connected. • Verify the IP address and DNS server settings. • Verify if the firewall is configured to allow outbound traffic. • Verify if the ISP network is operating normally.
	Proxy settings are incorrect.	Set proxy host, port, username, and password as required.
	Message Forwarding Service is temporarily unavailable.	Wait for the service to become available.
The Message Forwarding Service is unable to connect to a device platform notification service. [Code 100-105, 200-202, 211-212]	The platform provider service is temporarily unavailable to the Message Forwarding Service.	Wait for the service to become available.
The device communication token is no longer registered with the platform provider service. [Code 203]	The OpenManage Mobile application has been updated, restored, uninstalled, or the device operating system has been upgraded or restored.	Reinstall OpenManage Mobile on the device or follow the OpenManage Mobile troubleshooting procedures specified in the <i>OpenManage Mobile User's</i>

Table 30. Troubleshooting OpenManage Mobile (continued)

Problem	Reason	Resolution
		<p>Guide and reconnect the device to OpenManage Enterprise.</p> <p>If the device is no longer connected to OpenManage Enterprise, remove the subscriber.</p>
The OpenManage Enterprise registration is being rejected by the Message Forwarding Service. [Code 154]	An obsolete version of OpenManage Enterprise is being used.	Upgrade to a newer version of OpenManage Enterprise.

Related tasks

OpenManage Mobile settings on page 185

Related information

OpenManage Mobile settings on page 185

Other references and field descriptions

Definitions about some of the commonly displayed fields on the OpenManage Enterprise Graphical User Interface (GUI) are listed and defined in this chapter. Also, other information that is useful for further reference is described here.

Topics:

- Firmware and DSU requirement for HTTPS
- Schedule Reference
- Firmware baseline field definitions
- Supported and unsupported actions on 'Proxied' sleds
- Schedule job field definitions
- Alert categories after EEMI relocation
- Token substitution in remote scripts and alert policy
- Field service debug workflow
- Unblock the FSD capability
- Install or grant a signed FSD DAT.ini file
- Invoke FSD
- Disable FSD
- Catalog Management field definitions
- Firmware/driver compliance baseline reports— devices with 'Unknown' compliance status
- Generic naming convention for Dell EMC PowerEdge servers

Firmware and DSU requirement for HTTPS

If you have enabled the use of HTTPS for network share operations, then the servers must have the following minimum firmware and DSU to support the HTTPS-enabled device operations:

Use Case / Operation	YX2X (12G) or YX3X (13G) servers	YX4X (14G) and above servers
Firmware Update	FW v. 2.70.70.70	FW v. 3.00.00.00
Driver Update	DSU v.1.9.1	DSU v.1.9.1
Server Configuration Profile (SCP) for template capture, deployment, configuration inventory, and remediation)	FW v. 2.70.70.70	FW v. 3.00.00.00
Technical Support Report (TSR)	N/A	FW v. 3.21.21.21
Remote Diagnostics	N/A	FW v. 3.00.00.00

Schedule Reference

- **Update Now:** The firmware version is updated and matched to the version available in the associated catalog. To make the update become effective during the next device restart, select the **Stage for next server reboot** check box.
- **Schedule Later:** Select to specify a date and time when the firmware version must be updated.

Firmware baseline field definitions

- **COMPLIANCE:** The health status of the firmware baseline. Even if one device associated with a firmware baseline is in critical health status, the baseline health itself is declared as critical. This is called the rollup health status, which is equal to the status of the baseline that has high severity. For more information about Rollup Health status, see the *MANAGING THE ROLLUP HEALTH STATUS BY USING IDRAC ON THE DELL EMC 14TH GENERATION AND LATER POWEREDGE SERVERS* white paper on the Dell TechCenter.
- **NAME:** The firmware baseline name. Click to view the baseline compliance report on the **Compliance Report** page. For more information about creating a firmware baseline, see *Create a firmware/driver baseline* on page 84.
- **CATALOG:** The firmware catalog to which the firmware baseline belongs to. See *Manage firmware and driver Catalogs* on page 81.
- **LAST RUN TIME:** The time when the baseline compliance report is last run. See *Check the compliance of a device firmware and driver* on page 86.

Supported and unsupported actions on 'Proxied' sleds

Some device actions are not available for sleds in a 'Proxied' Managed State. The following table shows supported and unsupported Redfish- actions on the proxied sleds.

Capability_ID	Action	Action_Description	RedFish
1	POWER_CONTROL_ON	Power up	YES
2	POWER_CONTROL_OFF	Power Down hard/graceful	YES
3	POWER_CONTROL_RESET	Power reset hard/graceful	YES
4	SENSOR_DETAILS	Get Sensor Info, sub system health details	No
5	POWER_MONITOR	Power statistics retrieval	YES
6	TEMPERATURE_MONITOR	Temp statistics retrieval	YES
8	FW_UPDATE	Remote Firmware update capability.	YES
9	BLINK_LED	Identify function on server	YES
11	HW_LOGS	System Hardware logs	YES
12	DIAGS	Diagnostics	No
13	TSR	Tech Support Report	No
16	VIRTUAL_CONSOLE	Ability to execute RACADM tasks	No
30	REMOTE_RACADM	14G specific features	No
31	REMOTE_IPMI	14G specific features	No
32	REMOTE_SSH	14G specific features	No

Schedule job field definitions

- **Run now** to start the job immediately.
- **Run Later** to specify a later date and time.

- **Run On Schedule** to run repeatedly based on a selected frequency. Select **Daily**, and then select the frequency appropriately.

NOTE: By default, the job scheduler clock is reset at 12:00 A.M. everyday. The cron format does not consider the job creation time while calculating the job frequency. For example, if a job is started at 10:00 A.M. to run after every 10 hours, the next time the job runs is at 08:00 P.M. However, the subsequent time is not 06:00 A.M. next day but 12:00 A.M. This is because the scheduler clock is reset at 12:00 A.M. everyday.

Alert categories after EEMI relocation

Table of EEMI relocations

Table 31. Alert categories in OpenManage Enterprise

Previous Category	Previous Subcategory	New Category	New Subcategory
Audit	Devices	System Health	Devices
Audit	Devices	Configuration	Devices
Audit	Devices	Configuration	Devices
Audit	Devices	Configuration	Devices
Audit	Devices	Configuration	Devices
Audit	Application	Configuration	Application
Audit	Application	Configuration	Application
Audit	Application	Configuration	Application
Audit	Application	Configuration	Application
Audit	Devices	Audit	Users
Audit	Templates	Configuration	Templates
Audit	Templates	Configuration	Templates
Audit	Templates	Configuration	Templates
Audit	Templates	Configuration	Templates
Audit	Templates	Configuration	Templates
Configuration	Inventory	Configuration	Job
Configuration	Inventory	Configuration	Job
Configuration	Inventory	Configuration	Job
Configuration	Inventory	Configuration	Devices
Configuration	Inventory	Configuration	Devices
Configuration	Inventory	Configuration	Devices
Configuration	Firmware	Configuration	Jobs
Configuration	Firmware	Configuration	Jobs
Miscellaneous	Jobs	Configuration	Jobs
Miscellaneous	Jobs	Configuration	Jobs
Miscellaneous	Jobs	Configuration	Jobs
Miscellaneous	Generic	Configuration	Generic
Miscellaneous	Generic	Configuration	Generic

Table 31. Alert categories in OpenManage Enterprise (continued)

Previous Category	Previous Subcategory	New Category	New Subcategory
Miscellaneous	Generic	Configuration	Generic
Miscellaneous	Generic	Configuration	Generic
Miscellaneous	Generic	Configuration	Generic
Miscellaneous	Generic	Configuration	Generic
Miscellaneous	Generic	Configuration	Generic
Miscellaneous	Generic	Configuration	Generic
Miscellaneous	Devices	Configuration	Devices
Miscellaneous	Devices	Configuration	Devices
Audit	Security	Configuration	Security
Audit	Security	Configuration	Security
Audit	Security	Configuration	Security

Token substitution in remote scripts and alert policy

OpenManage Enterprise supports use of tokens to enhance remote scripting and creation of the alert policies.

Table 32. Tokens supported in OpenManage Enterprise

Tokens	Description
\$IP	Device IP Address
\$MSG	Message
\$DATE	Date
\$TIME	Time
\$SEVERITY	Severity
\$SERVICETAG	Service tag
\$RESOLUTION	Recommended Resolution
\$CATEGORY	Alert Category Name
\$ASSETTAG	Asset tag
\$MODEL	Model Name
\$HOSTNAME	FQDN or Hostname (if FQDN is not present)

Field service debug workflow

In OpenManage Enterprise, you can authorize console debugging by using the Field Service Debug (FSD) option.

About this task

By using FSD, you can perform the following tasks:

- Allow enabling and copying of debug logs
- Allow copying of real-time logs
- Allow backing up or restoring of database to VM.

The topics referenced in each task provide detailed instructions. To enable FSD, perform the following tasks:

Steps

1. Unblock FSD capability. See [Unblock the FSD capability](#) on page 195.
2. Install or grant signed FSD DAT.ini file. See [Install or grant a signed FSD DAT.ini file](#) on page 195.
3. Invoke FSD. See [Invoke FSD](#) on page 196.
4. Disable FSD. See [Disable FSD](#) on page 196.


Unblock the FSD capability

About this task

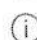
You can unblock the FSD capability through the TUI screen.

Steps

1. Navigate to the TUI main menu.
2. On the TUI screen, to use the FSD option, select **Enable Field Service Debug (FSD) Mode**.
3. To generate a new FSD unblock request, on the **FSD Functions** screen, select **Unblock FSD Capabilities**.
4. To determine the duration of the debug capabilities being requested, select a start and end date.
5. On the **Choose Requested Debug Capabilities** screen, select a debug capability from a list of debug capabilities unique to the console. In the lower-right corner, select **Generate**.

 **NOTE:** The debug capability that is currently supported is `RootShell`.

- a. You can download the generated .dat file from the **Audit Logs > Troubleshoot** menu in UI. Then, to complete the FSD enablement, upload the signed .dat file and SSH public key using the Upload options under the **Troubleshoot** menu.
 - b. If you have enabled CIFS share using the **Enable CIFS share for FSD (emergency use only)** option in TUI, then, use **Install/Grant Signed FSD DAT File** on the FSD Functions screen. See [Install or grant a signed FSD DAT.ini file](#) on page 195.
6. On the **Download DAT file** screen, view the signing instructions and the URL address of the share where the DAT.ini file exists.
 7. Use an external client to extract the DAT.ini file from the URL address of the share mentioned in step 6.

 **NOTE:**

- The download share directory has read-only privileges and supports only one DAT.ini file at a time.
- If the DAT file is downloaded as DAT.txt, you must rename it to DAT.ini.

8. Perform either of the following tasks depending on whether you are an external user or an internal Dell EMC user:
 - Send the DAT.ini file to a Dell EMC contact for signing if you are an external user.
 - Upload the DAT.ini file to appropriate Dell Field Service Debug Authentication Facility (FSDAF) and submit.
9. Wait for a Dell EMC signed and approved DAT.ini file to be returned.

Install or grant a signed FSD DAT.ini file

Prerequisites

Ensure that you have received the DAT.ini file, which is signed and approved by Dell EMC.

About this task

 **NOTE:**

- If the DAT file is downloaded as DAT.txt, you must rename it to DAT.ini.
- After Dell EMC approves the DAT.ini file, you must upload the file to the console appliance that generated the original unblock command.

Steps

1. To upload a signed DAT.ini file, on the **FSD Functions** screen, select **Install/Grant Signed FSD DAT File**.

NOTE: The upload share directory has write-only privileges and supports only one DAT.ini file at a time. The DAT.ini file size limit is 4 KB.

2. On the **Upload signed DAT file** screen, follow the instructions about uploading the DAT.ini file to a given file share URL.
3. Use an external client to upload the DAT.ini file to a share location.
4. On the **Upload signed DAT file** screen, select **I have uploaded the FSD DAT file**.

Results

If there are no errors during DAT.ini file upload, a message confirming the successful installation of the certificate is displayed. To continue, click **OK**.

The DAT.ini file upload can fail because of any of the following reasons:

- The upload share directory has insufficient disk space.
- The uploaded DAT.ini file does not correspond to the previous debug capability request.
- The signature provided by Dell EMC for the DAT.ini file is not valid.

Invoke FSD

Prerequisites

Ensure that the DAT.ini file is signed, returned by Dell EMC, and uploaded to OpenManage Enterprise.

Steps

1. To invoke a debug capability, on the **FSD Functions** screen, select **Invoke FSD Capabilities**.
2. On the **Invoke Requested Debug Capabilities** screen, select a debug capability from a list of debug capabilities that is approved in the Dell EMC signed DAT.ini file. In the lower-right corner, click **Invoke**.

NOTE: The debug capability that is currently supported is, RootShell.

Next steps

While the `invoke` command is run, OpenManage Enterprise can start an SSH daemon. The external SSH client can attach with OpenManage Enterprise for debugging purposes.

Disable FSD

About this task

After you invoke a debug capability on a console, it continues to operate until the console is restarted, or the debug capability is stopped. Else, the duration determined from the start and end date exceeds.

Steps

1. To stop the debug capabilities, on the **FSD Functions** screen, select **Disable Debug Capabilities**.
2. On the **Disable Invoked Debug Capabilities** screen, select a debug capability or capabilities from a list of currently invoked debug capabilities. From the lower right corner of the screen, select **Disable**.

Results

Ensure that you stop any SSH daemon or SSH sessions that are currently using the debug capability.

Catalog Management field definitions

CATALOG NAME: Name of the catalog. Built-in catalogs cannot be edited.

DOWNLOAD: Indicates the download status of catalogs from its repository folder. Statuses are: Completed, Running, and Failed.

REPOSITORY: Repository types such as Dell.com, CIFS, and NFS.

REPOSITORY LOCATION: Location where the catalogs are saved. Examples are Dell.com, CIFS, and NFS. Also, indicates the completion status of a job running on the catalog.

CATALOG FILE: Type of catalog file.

CREATED DATE: Date when the catalog file was created.

Firmware/driver compliance baseline reports— devices with 'Unknown' compliance status

The firmware or driver compliance status of the following storage, networking, and hyperconverged infrastructure (HCI) devices in the firmware/driver baseline compliance reports is displayed as Unknown as the Dell firmware/driver catalog does not support the firmware or software updates for these devices.

Table 33. Firmware/driver compliance baseline reports—'false' compliant devices

Device Category	Device List
Storage	<ul style="list-style-type: none">• SC Series• MD Series• ME Series
Network devices in the FX2, VRTX, and M1000e chassis	<ul style="list-style-type: none">• F10 switches• IOAs (Input/Output Aggregators)• IOMs (Input/Output Modules)
Hyperconverged Appliances (HCI)	<ul style="list-style-type: none">• VXRail• XC Series
Devices updatable using individual device's Dell Update Package (DUP) but not directly supported on Dell catalog	<ul style="list-style-type: none">• MX9116n Fabric Engine• MX5108n Ethernet Switch• PowerEdge MX5000s
Devices that cannot be updated using the Dell catalog or the individual DUP NOTE: For firmware/driver update of these devices, please refer the respective device's Installation Guide.	<ul style="list-style-type: none">• MX7116n Fabric Expander Module• PowerEdge MX 25GbE PTM

NOTE: For the complete list of devices in the SC, MD, ME, and XC series, refer https://topics-cdn.dell.com/pdf/dell-openmanage-enterprise_compatibility-matrix2_en-us.pdf

Generic naming convention for Dell EMC PowerEdge servers

To cover a range of server models, the PowerEdge servers are now be referred to using the generic naming convention and not their generation.

This topic explains how to identify the generation of a PowerEdge server that are referred to using the generic naming convention.

Example:

The R740 server model is a rack, two processor system from the 14th generation of servers with Intel processors. In the documentation, to refer to R740, generic naming convention **YX4X** server is used, where:

- The letter **Y** (alphabet) is used to denote the following server form factors:
 - **C** = Cloud - Modular server nodes for hyper-scale environments
 - **F** = Flexible - Hybrid rack-based sleds for rack-based FX2/FX2s enclosure
 - **M** or **MX*** = Modular - Blade servers for the modular enclosure MX7000, M1000e and/or VRTX
 - **R** = Rack-mountable servers

- **T** = Tower Servers
- The letter **X** (digit) denotes the class (number of processors) of the server.
- The digit **4** denotes the generation of the server.
- The letter **X** (digit) denotes the make of the processor.

Table 34. PowerEdge servers naming convention and examples

YX3X servers	YX4X systems
PowerEdge M630	PowerEdge M640
PowerEdge M830	PowerEdge R440
PowerEdge T130	PowerEdge R540



OpenManage Enterprise Solution Brief

Many IT organizations are dealing with increased workloads and decreased budgets. Under these conditions, IT professionals often have less time and fewer resources. It is critical for IT professionals to work smarter.


Dell EMC OpenManage Enterprise is an intuitive infrastructure management console. It is designed to take the complexity out of IT infrastructure management. It delivers better results with less time and fewer steps. OpenManage Enterprise helps IT professionals balance time and energy between complex IT infrastructure and business goals.

SIMPLIFY



Robust, intuitive, management capabilities, regardless of form-factor

UNIFY



One-to-many management from a single console: built for scale

AUTOMATE



Automated IT processes for greater efficiency

SECURE



Design for security throughout the infrastructure lifecycle

Dell EMC OpenManage Enterprise


Simplify: OpenManage Enterprise reduces learning time with HTML5 GUI with elastic search engine. It navigates to critical information and tasks easier and quicker. The automatable processes, templates and policies can be created and edited through a simple menu driven method.

Unify: OpenManage Enterprise supports up to 8,000 devices regardless of form factors. It supports Dell EMC PowerEdge racks, towers, and modular servers. It also monitors and creates alerts for third-party devices or PowerVault MD and ME Storage systems.

The innovative plugin design provides future extensibility. Through the same interface, plugins can be easily installed, updated and disabled. The first plugin will be OpenManage Enterprise – Power Manager.¹

Automate: From discovery to retirement, activities can be managed in the same console. In minutes, devices can be deployed automatically with templates based on service tags or node IDs.

Secure: Security is always the top priority. To protect your infrastructure, OpenManage Enterprise detects drift from a user-defined configuration template, alerts users, and remediates misconfigurations based on pre-setup policies.



Dell EMC OpenManage Enterprise systems management console

Advantages

Reduce the time and effort required to manage IT environments seamlessly with one unified tool

- Monitor third-party infrastructure in the same tool
- Require minimal training through an intuitive dashboard and elastic search engine
- Reduce repetitive tasks with templates
- Deploy as a secure virtual appliance, supporting ESXi, Hyper-V and KVM environments
- Integrate the OpenManage Mobile application to receive alerts anytime, anywhere

Simplify – robust, intuitive, management capabilities, regardless of form-factor

Features	Description	Benefits
Modern user interface with elastic search capabilities	Leverage the modern HTML5 standard while enabling an elastic search engine. Allow IT pros to find anything within the console in a single search	Minimize training time and maximize efficiency by delivering quick results on searches involving devices, hardware, and software inventory
Flexible mobile notifications	Integrate the OpenManage Mobile application to stay connected with your data center	Provide visible notification to the IT infrastructure and data center events anytime, anywhere
Full-lifecycle configuration management	Manage, deploy, and monitor server, chassis, and IOA through editable templates	Gain time back to focus on management tasks that drive more value for your business and your customers
Accurate log information	Locate root causes with detailed log information	Gain immediate insights for remediations

Unify – One-to-many management from a single console; built for scale


Features	Description	Benefits
Full infrastructure management	Manage up to 8,000 devices regardless of form factors - PowerEdge tower, rack, and modular. Monitor 3 rd party infrastructure	Reduce the time and effort required to manage and monitor IT environments seamlessly
Extendable plug-in architecture	Extend management capabilities with an intuitive plug-in architecture that integrates data center management tasks into a single interface	Streamline and enable power management from the intuitive OpenManage Enterprise interface
Extended modular support	Support modular servers, storage, and networking sleds with OpenManage Enterprise - Modular edition	Require minimal training through integration between OpenManage Enterprise and OpenManage Enterprise - Modular Edition

Automate – Automated IT processes for greater efficiency

Features	Description	Benefits
Streamlined remote management	Create a series of remote commands in a single batch, run immediately or schedule for later	Maximize IT efficiency and minimize IT downtime by automating a series of tasks
Automated server deployment	Automatically apply a template to selected devices based on service tag or node IDs	Decrease deployment time while preventing costly errors and downtime
Dynamic update repository refresh	Create or schedule searches for new available updates on Dell.com or through Dell Repository Manager. Maintain up-to-date repositories from OpenManage Enterprise interface	Maximize efficiency by identifying new available updates for systems or software within users' infrastructure
Customizable reports	Create customized reports to fit your business needs – for example, to quickly locate and filter NIC card information when there is a recall by vendor	Align automated reports to your business needs

Secure – Design for security


Features	Description	Benefits
Packaged as a virtual appliance	Readily deploy as a virtual appliance in ESXi, Hyper-V and KVM environments	High security standard throughout appliance testing, development, deployment, and user experience
Configuration and firmware drift detection	Create firmware and configuration baselines for compliance monitoring and enable automated updates on your schedule	Employ baselines to maintain security standards, performance optimizations and management conformity
Customizable alert policies	Build and design customized alert notifications that align with your business needs – detect, notify, and remediate	Improve efficiency and security by alerting the right contacts at the right time in the right way



Learn more about Dell EMC
OpenManage Enterprise



Contact a Dell EMC Expert



View more resources



Join the conversation
with #HashTag

Dell Technologies PowerEdge RAID Controller 11 User's Guide

PERC H755 adapter, H755 front SAS, H755N front NVMe, H755 MX adapter, H750 adapter SAS, H355 adapter SAS, H355 front SAS, and H350 adapter SAS

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
Operating systems supported by PERC 11 cards.....	12
Technical specifications of PERC 11 cards.....	13
 Chapter 2: Applications and User Interfaces supported by PERC 11.....	 16
Comprehensive Embedded Management	16
Dell OpenManage Storage Management.....	16
Human Interface Infrastructure Configuration Utility.....	16
The PERC Command Line Interface.....	17
 Chapter 3: Features of PowerEdge RAID Controller 11.....	 18
Controller features.....	18
Non-Volatile Memory Express.....	18
Opal Security Management.....	19
Hardware Root of Trust	19
1 MB I/O.....	19
Autoconfigure RAID 0.....	19
Disk roaming.....	20
FastPath.....	20
Non-RAID disks.....	21
Physical disk power management.....	21
Profile Management.....	21
Secure firmware update.....	21
Snapdump.....	21
Virtual disk features.....	21
Virtual disk write cache policy.....	22
Virtual disk read cache policy.....	22
Virtual disk migration.....	23
Virtual disk initialization.....	23
Full initialization.....	23
Fast initialization.....	23
Reconfiguration of virtual disks.....	24
Background operations.....	26
Background initialization.....	26
Consistency checks.....	26
Hard drive features.....	26

Self-Encrypting Disks.....	26
Instant secure erase.....	27
4 KB sector disk drives.....	27
Fault tolerance.....	27
The SMART feature.....	27
Patrol Read.....	28
Physical disk failure detection.....	28
Controller cache.....	29
Battery Transparent Learn Cycle.....	30
Linux operating system device enumeration.....	30
Chapter 4: Install and remove a PERC 11 card.....	32
Safety instructions.....	32
Before working inside your system.....	33
After working inside your system.....	33
Remove the PERC H755 adapter.....	33
Install the PERC H755 adapter.....	34
Remove the PERC H755 front SAS card.....	35
Install the PERC H755 front SAS card.....	36
Remove the PERC H755N front NVMe card.....	37
Install the PERC H755N front NVMe card.....	39
Remove the PERC H755 MX adapter.....	40
Install the PERC H755 MX adapter.....	41
Remove the PERC H750 adapter SAS.....	43
Install the PERC H750 adapter SAS.....	43
Remove the PERC H355 adapter SAS.....	44
Install the PERC H355 adapter SAS.....	45
Remove the PERC H355 front SAS.....	46
Install the PERC H355 front SAS card.....	48
Remove the PERC H350 adapter SAS.....	49
Install the PERC H350 adapter SAS.....	50
Chapter 5: Driver support for PERC 11	52
Creating the device driver media.....	52
Download and save PERC 11 drivers from the support site.....	52
Download and save PERC 11 drivers from the Dell Systems Service and Diagnostic Tools	52
Windows driver installation.....	53
Install PERC 11 driver while newly installing the Windows Server 2016 and later.....	53
Install PERC 11 driver on which the Windows Server 2016 is already installed and later.....	53
Update PERC 11 driver that runs on Windows Server 2016 and later.....	54
Linux driver installation.....	54
Install or update a RPM driver package using the KMOD support.....	55
Install or update a RPM driver package using the KMP support.....	55
Loading the driver while installing an operating system.....	56
Chapter 6: Firmware.....	57
Update firmware controller using Dell Update Package (DUP).....	57
Chapter 7: Manage PERC 11 controllers using HII configuration utility.....	58

Enter the PERC 11 HII configuration utility.....	58
Exit the PERC 11 HII configuration utility.....	58
Navigate to Dell PERC 11 configuration utility.....	59
View the HII Configuration utility dashboard.....	59
Configuration management.....	60
Auto Configure RAID 0.....	60
Create virtual disks.....	60
Create profile based virtual disk.....	61
View disk group properties.....	62
Convert to Non-RAID disk.....	62
Delete configurations.....	62
Controller management.....	62
Clear controller events.....	62
Save controller events.....	63
Save debug log.....	63
Enable security.....	63
Disable security.....	63
Change security settings.....	64
Restore factory default settings.....	64
Auto configure behavior.....	64
Manage controller profile.....	64
Advanced controller properties.....	65
Virtual disk management.....	68
Virtual disk numbering.....	68
Configure Virtual Disks.....	70
Perform expand virtual disk operation.....	70
Perform consistency check.....	70
Physical disk management.....	71
View physical disk properties.....	71
Cryptographic erase.....	72
Physical disk erase.....	72
Assigning a global hot spare.....	73
Assigning a dedicated hot spare.....	73
Convert to Non-RAID disk.....	74
Hardware components.....	74
View battery properties.....	74
View physical disks associated with an enclosure.....	75
Security key management in HII configuration utility.....	75
Chapter 8: Security key and RAID management.....	77
Security key implementation.....	77
Local Key Management.....	77
Create a security key.....	77
Change Security Settings.....	78
Disable security key.....	78
Create a secured virtual disk.....	79
Secure a non-RAID disk.....	79
Secure a pre-existing virtual disk.....	79
Import a secured non-RAID disk.....	79
Import a secured virtual disk.....	80

Dell Technologies OpenManage Secure Enterprise Key Manager.....	80
Supported controllers for OpenManage Secure Enterprise Key Manager.....	80
Manage enterprise key manager mode.....	81
Disable enterprise key manager mode.....	81
Manage virtual disks in enterprise key manager mode.....	81
Manage non-RAID disks in enterprise key manager mode.....	81
Transition of drives from local key management to enterprise key management (with supported firmware for PERC and iDRAC).....	81
Transition of drives from local key management to enterprise key management (without supported firmware for PERC and iDRAC).....	82
Chapter 9: Troubleshooting.....	83
Single virtual disk performance or latency in hypervisor configurations.....	83
Configured disks removed or not accessible error message.....	83
Dirty cache data error message.....	84
Discovery error message.....	84
Drive Configuration Changes Error Message.....	84
Windows operating system installation errors	85
Firmware fault state error message.....	85
Foreign configuration found error message.....	85
Foreign configuration not found in HII error message.....	85
Degraded state of virtual disks.....	85
Memory errors.....	85
Preserved Cache State.....	86
Security key errors.....	86
Secured foreign import errors.....	86
Failure to select or configure non Self-Encrypting Disks non-SED.....	86
Failure to delete security key.....	86
Failure of Cryptographic Erase on encryption-capable physical disks.....	87
General issues.....	87
PERC card has yellow bang in Windows operating system device manager.....	87
PERC card not seen in operating systems.....	87
Physical disk issues.....	87
Physical disk in failed state.....	87
Unable to rebuild a fault tolerant virtual disk.....	87
Fatal error or data corruption reported.....	87
Multiple disks are inaccessible.....	88
Rebuilding data for a failed physical disk.....	88
Virtual disk fails during rebuild using a global hot spare.....	88
Dedicated hot spare disk fails during rebuild.....	88
Redundant virtual disk fails during reconstruction.....	89
Virtual disk fails rebuild using a dedicated hot spare.....	89
Physical disk takes a long time to rebuild.....	89
Drive removal and insertion in the same slot generates a foreign configuration event	89
SMART errors.....	89
Smart error detected on a non-RAID disk.....	89
Smart error detected on a physical disk in a non-redundant virtual disk.....	90
Smart error detected on a physical disk in a redundant virtual disk.....	90
Replace member errors.....	90
Source disk fails during replace member operation.....	90

Target disk fails during replace member operation.....	90
A member disk failure is reported in the virtual disk which undergoes replace member operation.....	91
Linux operating system errors.....	91
Virtual disk policy is assumed as write-through error message.....	91
Unable to register SCSI device error message.....	91
Drive indicator codes.....	92
HII error messages.....	92
Unhealthy Status of the Drivers.....	92
Rebuilding a drive during full initialization.....	93
System reports more drive slots than what is available.....	93
Chapter 10: Appendix RAID description.....	94
Summary of RAID levels.....	94
RAID 10 configuration.....	95
RAID terminology.....	96
Disk striping.....	96
Disk mirroring.....	96
Spanned RAID levels.....	97
Parity data.....	97
Chapter 11: Getting help.....	98
Recycling or End-of-Life service information.....	98
Contacting Dell.....	98
Locating the Express Service Code and Service Tag.....	98
Receiving automated support with SupportAssist	99
Chapter 12: Documentation resources.....	100

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, and PERC H755N front NVMe, PERC H750 adapter SAS, PERC H755 MX adapter, PERC H355 adapter SAS, PERC H355 front SAS, and PERC H350 adapter SAS cards which 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 Drive (SSD), and PCIe SSD (NVMe)
- Supported drive speeds for NVMe drives are 8 GT/s and 16 GT/s at maximum x2 lane width.

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

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

NOTE: PERC H750 adapter SAS, PERC H355 front SAS, PERC H355 adapter SAS, and PERC H350 adapter 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, and PERC H350 adapter SAS.

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
- Operating systems supported by PERC 11 cards
- Technical specifications of PERC 11 cards

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




Figure 2. Features of PERC H755 front SAS

- | | |
|------------------------------|--------------------------|
| 1. Battery | 2. PCIe input connector |
| 3. Power card edge connector | 4. Heatsink |
| 5. Backplane connector A | 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 |

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.

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

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 Gbps SAS	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	<ul style="list-style-type: none"> Without SAS Expander: 16 drives per controller With SAS Expander: Limited by platform offerings 	<ul style="list-style-type: none"> Without SAS Expander: 16 drives per controller With SAS Expander: Limited by platform offerings 	Not applicable	Limited by platform: 8 drives per controller	<ul style="list-style-type: none"> Without SAS Expander: 8 drives per controller With SAS Expander: Limited by platform offerings
NVMe maximum drive support	<ul style="list-style-type: none"> Without PCIe Switch Expander: 8 drives per controller With PCIe Switch Expander: Limited by platform offerings 	Not applicable	<ul style="list-style-type: none"> 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

- (i) **NOTE:** PERC H755 adapter and PERC H755 MX supports either SAS, SATA, or NVMe drives depending on the backplane/server configuration.
- (i) **NOTE:** PERC controller supports only conventional magnetic recording (CMR) drives, and does not support shingled magnetic recording (SMR) drives.
- (i) **NOTE:** PERC H755 family of controllers currently support SEKM starting with firmware version 52.14.0-3901.
- (i) **NOTE:** For information on number of drives in a disk group per virtual disk, see [Summary of RAID levels](#)
- (i) **NOTE:** As 14G PowerEdge Servers do not support Gen 4 speeds, PERC H750 adapter SAS will downtrain to Gen 3 speeds.

Feature	PERC H355 adapter SAS	PERC H355 front SAS	PERC H350 adapter SAS
RAID levels	0, 1, 10	0, 1, 10	0, 1, 10
Non-RAID	Yes	Yes	Yes
Enclosures per port	Not applicable	Not applicable	Not applicable
Processor	Broadcom RAID-onchip, SAS3816 chipset	Broadcom RAID-onchip, SAS3816 chipset	Broadcom RAID-onchip, SAS3816 chipset
Battery backup unit	No	No	No
Local Key Management security	No	No	No
Controller queue depth	1536	1536	1536

Feature	PERC H355 adapter SAS	PERC H355 front SAS	PERC H350 adapter SAS
Secure enterprise key manager security	No	No	No
Non-volatile cache	No	No	No
Cache memory	Not applicable	Not applicable	Not applicable
Cache function	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
Max no of disk groups	32	32	32
Max no of VDs per disk group	16	16	16
Hot swap devices supported	Yes	Yes	Yes
Autoconfig	Yes	Yes	Yes
Hardware XOR engine	Yes	Yes	Yes
Online capacity expansion	Yes	Yes	Yes
Dedicated and global hot spare	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)
VD strip size	64 KB	64 KB	64 KB
PCIe support	Gen 4	Gen 4	Gen 4
SAS/SATA maximum drive support	<ul style="list-style-type: none"> Without SAS Expander: 16 With SAS Expander: Limited by platform offering 	<ul style="list-style-type: none"> Without SAS Expander: 16 With SAS Expander: Limited by platform offering 	<ul style="list-style-type: none"> Without SAS Expander: 8 drives per controller With SAS Expander: Limited by platform offering
NVMe maximum drive support	Not applicable	Not applicable	Not applicable

① **NOTE:** As 14G PowerEdge Servers do not support Gen 4 speeds, PERC H350 adapter SAS will downtrain to Gen 3 speeds.

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 VDIs simultaneously with CEM.

Dell OpenManage Storage Management

Dell OpenManage Storage Management is a storage management application for Dell systems that provides enhanced features for configuring locally attached RAID disk storage. The Dell OpenManage storage management application enables you to perform controller and enclosure functions for all supported RAID controllers and enclosures from a single graphical or Command Line Interface (CLI). The User Interface (UI) is wizard-driven with features for novice and advanced users, and detailed online help. Using the Dell OpenManage storage management application, you can protect your data by configuring data-redundancy, assigning hot spares, or rebuilding failed physical disks. The fully featured CLI, which is available on select operating systems, allows you to perform RAID management tasks either directly from the console or through scripting.

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


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 EMC PowerEdge RAID Controller CLI Reference Guide* at www.dell.com/storagecontrollermanuals.

Features of PowerEdge RAID Controller 11

Topics:

- Controller features
- Virtual disk features
- Virtual disk initialization
- Reconfiguration of 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 physical disk 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, and PERC H350 adapter SAS.

Autoconfigure behavior

The autoconfigure behavior automatically configures unconfigured drives during reboot and hot insertion. As per the setting, unconfigured drives are configured as per the option; but the configured drives remain unaffected. PERC 11 supports **Off** and **Non-RAID** settings.

Table 2. 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 will remain unaffected
Off to Non-RAID disk	Unconfigured drives are converted to non-RAID disk; all the configured drives will remain unaffected
Non-RAID disk to Off	Unconfigured drives remain unconfigured good; all the configured drives will remain unaffected

NOTE: PERC H355 front SAS, PERC H355 adapter SAS, and PERC H350 adapter 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 www.dell.com/openmanagemanuals.

Profile Management

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

Table 3. 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, and PERC H350 adapter 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 EMC PowerEdge RAID Controller CLI Reference Guide](#) at www.dell.com/storagecontrollermanuals.

Virtual disk features

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

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

- Reconfiguration of virtual disk
- 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 4. 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. <i>i</i> 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. <i>i</i> NOTE: Certain data patterns and configurations perform better with a write-through cache policy.

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

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

Conditions under which write-back is 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 5. 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.

i **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 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.

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

i **NOTE:** Importing non-RAID drives and uneven span RAID 10 virtual disks from PERC 9 to PERC 11 is not supported.

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

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

i **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 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](#).

i **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](#).

i **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.

- ① **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.

Reconfiguration of virtual disks

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

- ① **NOTE:** Spanned virtual disks such as RAID 50 and 60 cannot be reconfigured.
- ① **NOTE:** Reconfiguring virtual disks typically impacts disk performance until the reconfiguration operation is complete.

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

1. If there is a single virtual disk in a disk group and free space is available, the capacity of a virtual disk can be expanded within that free space. If multiple virtual disks exist within a common disk group, the capacities of those virtual disks cannot be expanded.
① **NOTE:** Online capacity expansion is allowed on a disk group with a single virtual disk that begins at the start of the physical disk. It is not allowed when there is a free space at the beginning of a disk.
2. Free space is also available when the physical disks of a disk group are replaced by larger disks using the replace member feature. The capacity of a virtual disk can also be expanded by performing an OCE operation to add more physical disks.

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

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

- ① **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.
- ① **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.
- ① **NOTE:** The controller changes the write cache policy of all virtual disks to write-through until the RLM or OCE operation is complete.

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 6. RAID level migration

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

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

Background operations

Background initialization

Background initialization (BGI) is an automated process that writes parity or mirror data on newly created virtual disks. BGI does not run on RAID 0 virtual disks. You can control the BGI rate in the Dell OpenManage storage management application. Any change to the BGI rate does not take effect until the next BGI is executed.

i 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, and PERC H350 adapter 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](#).

i 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

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

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

SEDs that are secured by a non-PERC entity cannot be used by PERC. Ensure that the SED is re-provisioned in an applicable manner by that non-PERC entity before connecting to PERC.

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

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

i NOTE: PERC 11 supports Trusted Computing Group Enterprise (TCG) Security Subsystem Classes (SSC) SAS or SATA SED drives and TCG Opal SSC NVMe drives.