8

Security key and RAID management

Topics:

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.

Local Key Management

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

ⓘ **NOTE:** Under LKM, you are prompted for a passphrase when you create the key. This mode is not supported on PERC H355 adapter SAS and PERC H350 adapter 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 on supported platforms, see www.dell.com/idracmanuals.

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

Manage enterprise key manager mode

iDRAC manages Enterprise key manager features. For instructions on enabling enterprise key manager mode, see www.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. For more information, see Transition of drives from local key management to enterprise key (with supported firmware for PERC and iDRAC).

ⓘ **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 (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 https://www.dell.com/idracmanuals.

About this task

ⓘ **NOTE:** This feature is supported on firmware starting with version 52.16.1-4074.

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.

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 https://www.dell.com/supportassist.

9

Troubleshooting

To get help with your Dell Technologies PowerEdge RAID Controller 11 series, you can contact your Dell Technical Service representative or see https://www.dell.com/support.

Topics:

- Single virtual disk performance or latency in hypervisor configurations
- 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 error message
- Degraded state of virtual disks
- Memory errors
- Preserved Cache State
- Security key errors
- General issues
- Physical disk issues
- SMART errors
- Replace member errors
- Linux operating system errors
- Drive indicator codes
- HII error messages
- System reports more drive slots than what is available

Single virtual disk performance or latency in hypervisor configurations

Multi-initiator or hypervisor configurations running multiple I/O workloads to a single raid array may experience degraded performance or latency. This is caused by upper layers sending separate I/O workloads for each virtual machine to the storage subsystem which ends up being a random I/O workload to the 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

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: Turn off the system and reseat the controller.

For more information, see Install and remove a PERC 11 card.

Physical disk issues

Physical disk in failed state

Issue: One of the physical disks in the disk array is in the failed state.

Corrective Action: Update the PERC cards to the latest firmware available on https://www.dell.com/support and replace the drive.

Unable to rebuild a fault tolerant virtual disk

Issue: Cannot rebuild a fault tolerant virtual disk. For more information, see the alert log for virtual disks.

Probable Cause: The replacement disk is too small or not compatible with the virtual disk.

Corrective Action: Replace the failed disk with a compatible good physical disk with equal or greater capacity.

Fatal error or data corruption reported

Issue: Fatal error(s) or data corruption(s) are reported when accessing virtual disks.

Corrective Action: Contact 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 where to find reports of SMART errors that could indicate hardware failure, see the Dell OpenManage storage management documentation at www.dell.com/openmanagemanuals.

Smart error detected on a 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 message

Error:

```
<Date:Time> <HostName> kernel: sdb: asking for cache data failed<Date:Time>
<HostName> kernel: sdb: assuming drive cache: write through
```

Corrective Action: The error message is displayed when the Linux Small Computer System Interface (SCSI) mid-layer asks for physical disk cache settings. The controller firmware manages the virtual disk cache settings on a per controller and a per virtual disk basis, so the firmware does not respond to this command. The Linux SCSI mid-layer assumes that the virtual disk's cache policy is **Write-Through**. SDB is the device node for a virtual disk. This value changes for each virtual disk.

For more information about **Write-Through** cache, see Virtual Disk Write Cache Policies.

Except for this message, there is no effect of this behavior on normal operation. The cache policy of the virtual disk and the I/O throughput are not affected by this message. The cache policy settings for the PERC SAS RAID system remain unchanged.

Unable to register SCSI device error message

Error:

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

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

Drive indicator codes

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

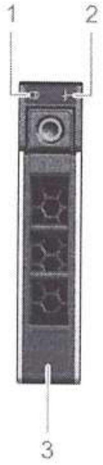

Figure 25. Drive indicators

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 18. 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 **BP15G+0.0** 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).

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

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 19. 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 20. 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 20. 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.

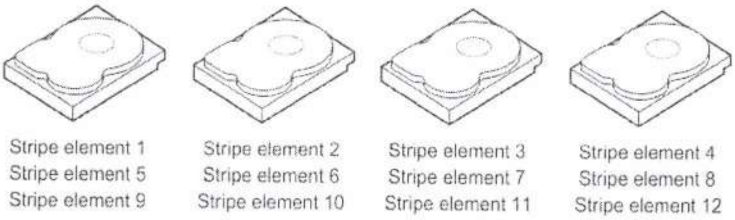

Figure 26. Example of disk striping (RAID 0)

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.

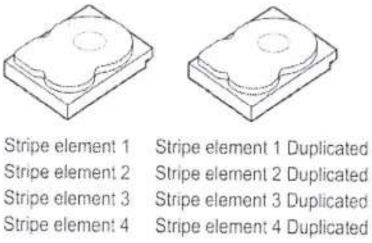
Stripe element 1
Stripe element 2
Stripe element 3
Stripe element 4
Stripe element 1 Duplicated
Stripe element 2 Duplicated
Stripe element 3 Duplicated
Stripe element 4 Duplicated

Figure 27. 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.

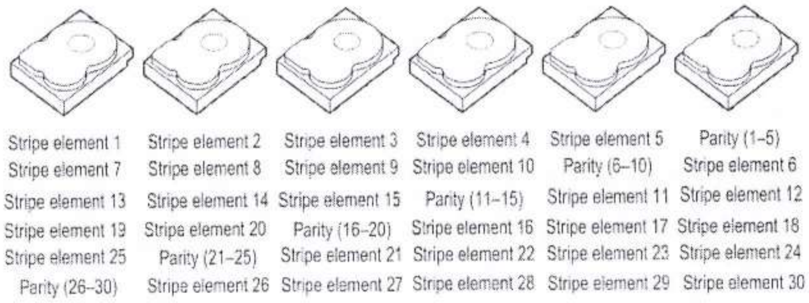

Stripe element 1	Stripe element 2	Stripe element 3	Stripe element 4	Stripe element 5	Parity (1–5)
Stripe element 7	Stripe element 8	Stripe element 9	Stripe element 10	Parity (6–10)	Stripe element 6
Stripe element 13	Stripe element 14	Stripe element 15	Parity (11–15)	Stripe element 11	Stripe element 12
Stripe element 19	Stripe element 20	Parity (16–20)	Stripe element 16	Stripe element 17	Stripe element 18
Stripe element 25	Parity (21–25)	Stripe element 21	Stripe element 22	Stripe element 23	Stripe element 24
Parity (26–30)	Stripe element 26	Stripe element 27	Stripe element 28	Stripe element 29	Stripe element 30

Figure 28. Example of Distributed Parity (RAID 5)

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

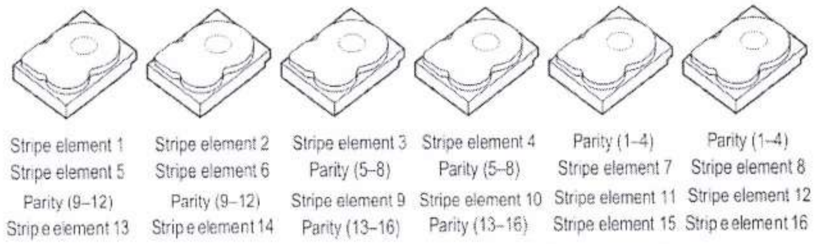

Stripe element 1	Stripe element 2	Stripe element 3	Stripe element 4	Parity (1–4)	Parity (1–4)
Stripe element 5	Stripe element 6	Parity (5–8)	Parity (5–8)	Stripe element 7	Stripe element 8
Parity (9–12)	Parity (9–12)	Stripe element 9	Stripe element 10	Stripe element 11	Stripe element 12
Stripe element 13	Stripe element 14	Parity (13–16)	Parity (13–16)	Stripe element 15	Stripe element 16

Figure 29. 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

Take back and recycling services are offered for this product in certain countries. If you want to dispose of system components, visit www.dell.com/recyclingworldwide and select the relevant country.

Contacting Dell

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

Steps

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

Locating the Express Service Code and Service Tag

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

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

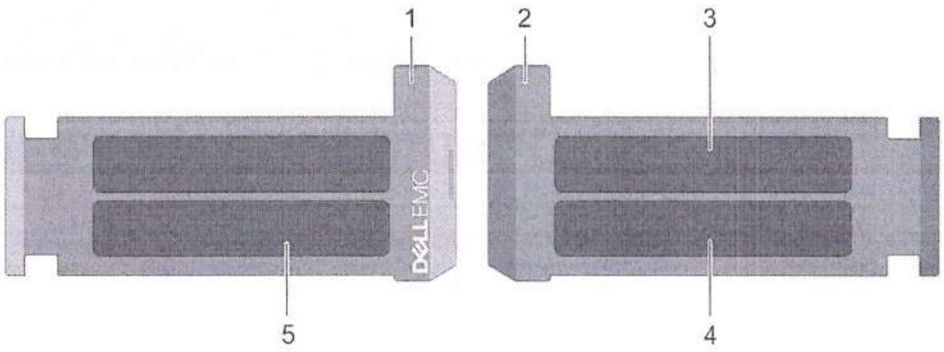

Figure 30. Locating the Express Service Code and Service tag

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

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

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

Receiving automated support with SupportAssist

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

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

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

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

     ⓘ NOTE: To locate the product name and model, see the front of your system.
  3. On the Product Support page, click **Manuals & documents**.
- Using search engines:
  - Type the name and version of the document in the search box.

Table 21. 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.	www.dell.com/poweredgemanuals
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.	www.dell.com/poweredgemanuals
	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 **?** > **About**.	www.dell.com/idracmanuals

Table 21. Additional documentation resources for your system (continued)

Task	Document	Location
	For information about installing the operating system, see the operating system documentation.	www.dell.com/operatingsystemmanuals
	For information about updating drivers and firmware, see the Methods to download firmware and drivers section in this document.	www.dell.com/support/drivers
Managing your system	For information about systems management software offered by Dell, see the Dell OpenManage Systems Management Overview Guide.	www.dell.com/poweredgemanuals
	For information about setting up, using, and troubleshooting OpenManage, see the Dell OpenManage Server Administrator User's Guide.	www.dell.com/openmanagemanuals > OpenManage Server Administrator
	For information about installing, using, and troubleshooting Dell OpenManage Enterprise, see the Dell OpenManage Enterprise User's Guide.	https://www.dell.com/openmanagemanuals
	For information about installing and using Dell SupportAssist, see the Dell EMC SupportAssist Enterprise User's Guide.	https://www.dell.com/serviceabilitytools
	For information about partner programs enterprise systems management, see the OpenManage Connections Enterprise Systems Management documents.	www.dell.com/openmanagemanuals
Understanding event and error messages	For information about the event and error messages generated by the system firmware and agents that monitor system components, go to qrl.dell.com > **Look Up** > **Error Code**, type the error code, and then click **Look it up**.	www.dell.com/qrl
Troubleshooting your system	For information about identifying and troubleshooting the PowerEdge server issues, see the Server Troubleshooting Guide.	www.dell.com/poweredgemanuals

Dell EMC PowerEdge R750

Technical Specifications

Regulatory Model: E70S Series
Regulatory Type: E70S001
July 2021
Rev. A01

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

1

Technical specifications

The technical and environmental specifications of your system are outlined in this section.

Topics:

Chassis dimensions

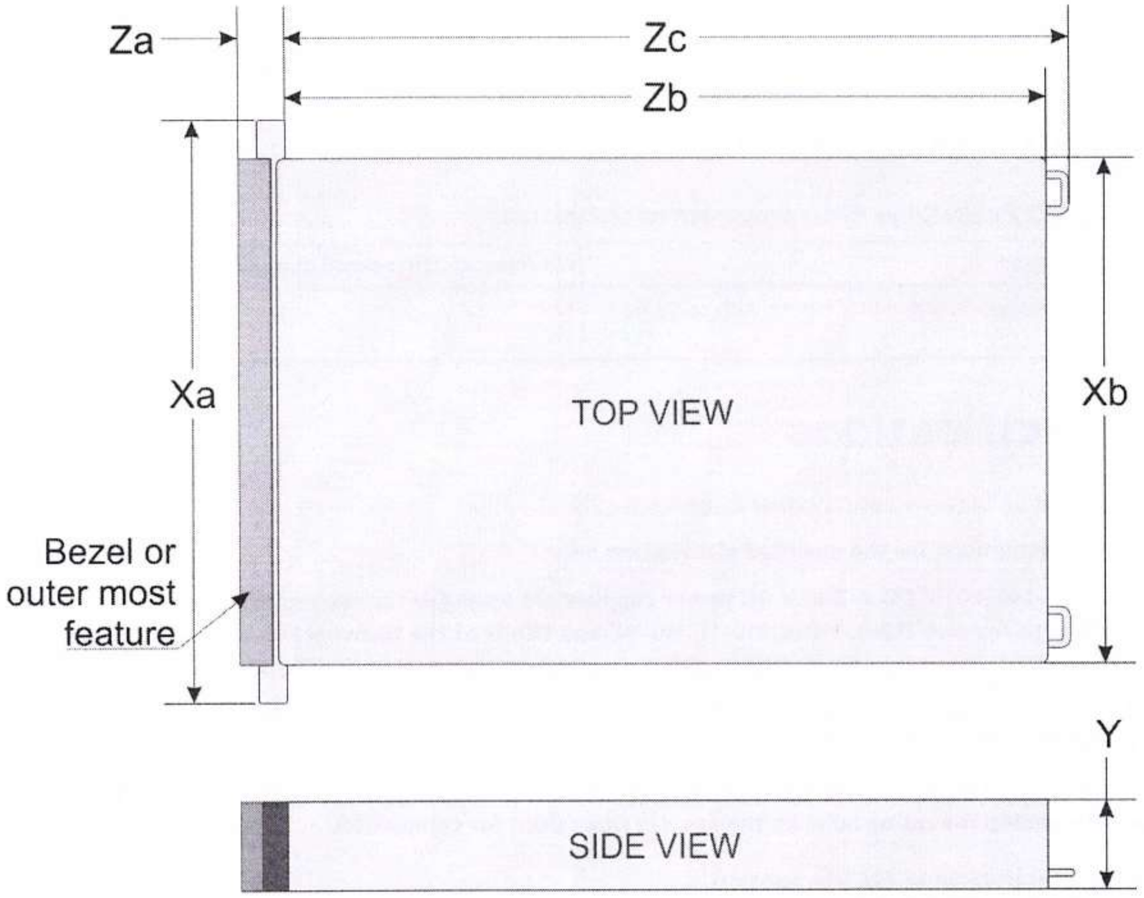

Figure 1. Chassis dimensions

Table 1. Chassis dimension for the system

Drives	Xa	Xb	Y	Za	Zb	Zc
0/8/12/16/24 drives	482.0 mm (18.97 inches)	434.0 mm (17.0 inches)	86.8 mm (3.41 inches)	35.84 mm (1.41 inches) with bezel 22.0 mm (0.86 inches) without bezel	700.7 mm (27.58 inches) Ear to rear wall	736.29 mm (28.92 inches) Ear to PSU handle

ⓘ **NOTE:** Zb is the nominal rear wall external surface where the system board I/O connectors reside.

Chassis weight

Table 2. Chassis weight

System configuration	Maximum weight (with all drives/SSDs)
0	27.7 kg (61.06 lb)
12 x 3.5-inch	35.3 kg (77.82 lb)
8 x 2.5-inch	29.6 kg (65.25 lb)

Table 2. Chassis weight (continued)

System configuration	Maximum weight (with all drives/SSDs)
16 x 2.5-inch	32.6 kg (71.87 lb)
24 x 2.5-inch	35.2 kg (77.60 lb)

Processor specifications

Table 3. Dell EMC PowerEdge R750 processor specifications

Supported processor	Number of processors supported
3rd Generation Intel Xeon Scalable processors with up to 40 cores	two

PSU specifications

The system supports up to two AC or DC power supply units (PSUs).

⚠ **WARNING: Instructions for the qualified electricians only:**

System using -(48-60) V DC or 240 V DC power supplies are intended for restricted access locations in accordance with Articles 110-5, 110-6, 110-11, 110-14, and 110-17 of the National Electrical Code, American National Standards Institute (ANSI)/National Fire Protection Association (NFPA) 70.

240 V DC power supplies shall be connected to the 240 V DC outlet from certified power distribution units if applicable in country or region of use.

Power supply cords/jumper cords and the associated plugs/inlets/connectors shall have appropriate electrical ratings referencing the rating label on the system when used for connection.

Table 4. PSU specifications for the system

PSU	Class	Heat dissipation (maximum)	Frequency	Voltage	Peak power	N/A	N/A	Peak power	N/A	Current
					High line/-72 VDC	High line/-72 VDC	High line/ 240 VDC	Low line/-40 VDC	Low line/-40 VDC	
800 W AC	Platinum	3139 BTU/hr	50/60 Hz	100 - 240 V	1360 W	800 W	800 W	1360 W	800 W	9.2 - 4.7 A
800 W Mixed Mode	N/A	3139 BTU/hr	N/A	240 V	1360 W	800 W	800 W	1360 W	800 W	3.8 A
1100 WDC	Titanium	4265 BTU/hr	N/A	-48 - -60 V	1870 W	1100 W	N/A	1870 W	1100 W	27.0 A
1100 W AC	Titanium	4299 BTU/hr	50/60 Hz	100 - 240 V	1870 W	1100 W	1100 W	1785 W	1050 W	12 - 6.3 A
1100 W Mixed Mode	N/A	4299 BTU/hr	N/A	240 V	1870 W	1100 W	1100 W	1870 W	1100 W	5.2 A
1400 W AC	Platinum	5459 BTU/hr	50/60 Hz	100 - 240 V	2380 W	1400 W	1400 W	1785 W	1050 W	12 - 8 A
1400 W Mixed Mode	N/A	5459 BTU/hr	N/A	240 V	2380 W	1400 W	1400 W	1785 W	1050 W	6.6 A