NetApp

Drives

SANtricity 11.6

NetApp
February 12, 2024

Table of Contents

Drives

Concepts

Drive terminology

Learn how the drive terms apply to your storage array.

Component	Description
DA	Data Assurance (DA) is a feature that checks for and corrects errors that might occur as data is transferred through the controllers down to the drives. Data Assurance can be enabled at the pool or volume group level, with hosts using a DA-capable I/O interface such as Fibre Channel.
Drive Security feature	Drive Security is a storage array feature that provides an extra layer of security with either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives. When these drives are used with the Drive Security feature, they require a security key for access to their data. When the drives are physically removed from the array, they cannot operate until they are installed in another array, at which point, they will be in a Security Locked state until the correct security key is provided.
Drive shelf	A drive shelf, also called an expansion shelf, contains a set of drives and two input/output modules (IOMs). The IOMs contain SAS ports that connect a drive shelf to a controller shelf or to other drive shelves.
DULBE	Deallocated or Unwritten Logical Block Error (DULBE) is an option on NVMe drives that allows a storage array to deallocate blocks that are part of a volume. Deallocating blocks on a drive can greatly reduce the time it takes to initialize volumes. In addition, hosts can deallocate logical blocks in the volume using the NVMe Dataset Management command.
FDE drives	Full Disk Encryption (FDE) drives perform encryption on the disk drive at the hardware level. The hard drive contains an ASIC chip that encrypts data during writes, and then decrypts data during reads.
FIPS drives	FIPS drives use Federal Information Processing Standards (FIPS) 140-2 level 2. They are essentially FDE drives that adhere to United States government standards for ensuring strong encryption algorithms and methods. FIPS drives have higher security standards than FDE drives.
HDD	Hard disk drives (HDDs) are data storage devices that use rotating metal platters with a magnetic coating.
Hot spare drives	Hot spares act as standby drives in RAID 1, RAID 5, or RAID 6 volume groups. They are fully functional drives that contain no data. If a drive fails in the volume group, the controller automatically reconstructs data from the failed drive to a hot spare.

Component	Description
NVMe	Non-Volatile Memory Express (NVMe) is an interface designed for flash-based storage devices, such as SSD drives. NVMe reduces I/O overhead and includes performance improvements, as compared to previous logical-device interfaces.
SAS	Serial Attached SCSI (SAS) is a point-to-point serial protocol that links controllers directly to disk drives.
Secure-capable drives	Secure-capable drives can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives, which encrypt data during writes and decrypt data during reads. These drives are considered secure-*capable* because they can be used for additional security using the Drive Security feature. If the Drive Security feature is enabled for volume groups and pools used with these drives, the drives become secure-*enabled*.
Secure-enabled drives	Secure-enabled drives are used with the Drive Security feature. When you enable the Drive Security feature and then apply Drive Security to a pool or volume group on secure-*capable* drives, the drives become secure-*enabled*. Read and write access is available only through a controller that is configured with the correct security key. This added security prevents unauthorized access to the data on a drive that is physically removed from the storage array.
SSD	Solid-state disks (SSDs) are data storage devices that use solid state memory (flash) to store data persistently. SSDs emulate conventional hard drives, and are available with the same interfaces that hard drives use.

Drive states

System Manager reports various states for drives.

Accessibility states

State	Definition
Bypassed	The drive is physically present, but the controller cannot communicate with it on either port.
Incompatible	One of the following conditions exists: • The drive is not certified for use in the storage array. • The drive has a different sector size. • The drive has unusable configuration data from an older or newer firmware version.
Removed	The drive has been improperly removed from the storage array.
Present	The controller can communicate with the drive on both ports.

State	Definition
Unresponsive	The drive is not responding to commands.

Role states

State	Definition
Assigned	The drive is a member of a pool or volume group.
In-use hot spare	The drive is currently being used as a replacement for a drive that has failed. Hot spares are used only in volume groups, not pools.
Standby hot spare	The drive is ready to be used as a replacement for a drive that has failed. Hot spares are used only in volume groups, not pools.
Unassigned	The drive is not a member of a pool or volume group.

Availability states

State	Definition
Failed	The drive is not working. The data on the drive is not available.
Impending Failure	It has been detected that the drive could fail soon. The data on the drive is still available.
Offline	The drive is not available for storing data usually because it is part of a volume group that is being exported or it is undergoing a firmware upgrade.
Optimal	The drive is working normally.

Solid State Disks (SSDs)

Solid-state disks (SSDs) are data storage devices that use solid state memory (flash) to store data persistently. SSDs emulate conventional hard drives, and are available with the same interfaces that hard drives use.

Advantages of SSDs

The advantages of SSDs over hard drives include:

- Faster start up (no spin up)
- Lower latency
- Higher I/O operations per second (IOPS)
- Higher reliability with fewer moving parts

- Lower power usage

- Less heat produced and less cooling required

Identifying SSDs

From the Hardware page, you can locate the SSDs in the front-shelf view. Look for drive bays that display a lightning bolt icon, which indicates an SSD is installed.

Volume groups

All drives in a volume group must be of the same media type (either all SSDs or all hard drives). Volume groups cannot have a mixture of media types or interface types.

Caching

The controller's write caching is always enabled for SSDs. Write caching improves performance and extends the life of the SSD.

In addition to the controller cache, you can implement the SSD cache feature to improve overall system performance. In SSD cache, the data is copied from volumes and stored on two internal RAID volumes (one per controller).

Hot spare drives

Hot spares act as standby drives in RAID 1, RAID 5, or RAID 6 volume groups for System Manager. They are fully functional drives that contain no data. If a drive fails in the volume group, the controller automatically reconstructs data from the failed drive to a drive assigned as a hot spare.

Hot spares are not dedicated to specific volume groups. They can be used for any failed drive in the storage array, as long as the hot spare and the drive share these attributes:

- Equal capacity (or greater capacity for the hot spare)

- Same media type (for example, HDD or SSD)

- Same interface type (for example, SAS)

How to identify hot spares

You can assign hot spares through the Initial Setup Wizard or from the Hardware page. To determine if hot spares are assigned, go to the Hardware page and look for any drive bays shown in pink.

How hot spare coverage works

Hot spare coverage works as follows:

- You reserve an unassigned drive as a hot spare for RAID 1, RAID 5, or RAID 6 volume groups.

  > ⓘ Hot spares cannot be used for pools, which have a different method of data protection. Instead of reserving an additional drive, pools reserve spare capacity (called *preservation capacity*) within each drive of the pool. If a drive fails in a pool, the controller reconstructs data in that spare capacity.

- If a drive within a RAID 1, RAID 5, or RAID 6 volume group fails, the controller automatically uses redundancy data to reconstruct the data from the failed drive. The hot spare is automatically substituted for the failed drive without requiring a physical swap.

- When you have physically replaced the failed drive, a copyback operation occurs from the hot spare drive to the replaced drive. If you have designated the hot spare drive as a permanent member of a volume group, the copyback operation is not needed.

- The availability of tray loss protection and drawer loss protection for a volume group depends on the location of the drives that comprise the volume group. The tray loss protection and drawer loss protection might be lost because of a failed drive and location of the hot spare drive. To make sure that tray loss protection and drawer loss protection are not affected, you must replace a failed drive to initiate the copyback process.

- The storage array volume remains online and accessible while you are replacing the failed drive, because the hot spare drive is automatically substituted for the failed drive.

Considerations for hot spare drive capacity

Select a drive with a capacity equal to or greater than the total capacity of the drive you want to protect. For example, if you have an 18-GiB drive with configured capacity of 8 GiB, you can use a 9-GiB or larger drive as a hot spare. Generally, do not assign a drive as a hot spare unless its capacity is equal to or greater than the capacity of the largest drive in the storage array.

> (i) If hot spares are not available that have the same physical capacity, a drive with lower capacity may be used as a hot spare if the "used capacity" of the drive is the same or smaller than the capacity of the hot spare drive.

Considerations for media and interface types

The drive used as a hot spare must share the same media type and interface type as the drives it will protect. For example, an HDD drive cannot serve as a hot spare for SSD drives.

Considerations for secure-capable drives

A secure-capable drive, such as FDE or FIPS, can serve as a hot spare for drives with or without security capabilities. However, a drive that is not secure-capable cannot serve as a hot spare for drives with security capabilities.

When you select a secure-enabled drive to be used for a hot spare, System Manager prompts you to perform a Secure Erase before you can proceed. The Secure Erase resets the drive's security attributes to secure-capable, but not secure-enabled.

> (i) When you enable the Drive Security feature and then create a pool or volume group from secure-capable drives, the drives become *secure-enabled*. Read and write access is available only through a controller that is configured with the correct security key. This added security prevents unauthorized access to the data on a drive that is physically removed from the storage array.

Recommended number of hot spare drives

If you used the Initial Setup wizard to automatically create hot spares, System Manager creates one hot spare for every 30 drives of a particular media type and interface type. Otherwise, you can manually create hot spare drives among the volume groups in the storage array.

How tos

Limit the drive view

If the storage array includes drives with different types of physical and logical attributes, the Hardware page provides filter fields that help you limit the drive view and locate specific drives.

About this task

The drive filters can limit the view to only certain types of physical drives (for example, all SAS), with certain security attributes (for example, secure-capable), at certain logical locations (for example, Volume Group 1). You can use these filters together or separately.

> (i) If all drives share the same physical attributes, the **Show drives that are…** filter field does not appear. If all drives share the same logical attributes, the **Anywhere in the storage array** filter field does not appear.

Steps

1. Select **Hardware**.

2. In the first filter field (under **Show drives that are…**), click the drop-down arrow to display the available drive types and security attributes.

   Drive types might include:

   - Drive media type (SSD, HDD)
   - Drive interface type (SAS, NVMe)
   - Drive capacity (highest to lowest)
   - Drive speed (highest to lowest) Security attributes might include:
   - Secure-capable
   - Secure-enabled
   - DA (Data Assurance) capable If any of these attributes are the same for all drives, they are not shown in the drop-down list. For example, if the storage array includes all SSD drives with SAS interfaces and speeds of 15000 RPM, but some SSDs have different capacities, the drop-down list displays only the capacities as a filtering choice.

   When you select an option from the field, the drives that do not match your filter criteria are grayed out in the graphic view.

3. In the second filter box, click the drop-down arrow to display the available logical locations for the drives.

   > (i) If you need to clear your filter criteria, select **Clear** on the far right of the filter boxes.

   Logical locations might include:

   - Pools
   - Volume Groups
   - Hot spare

- SSD Cache
- Unassigned When you select an option from the field, the drives that do not match your filter criteria are grayed out in the graphic view.

4. **Optional:** You can select **Turn on locator lights** at the far right of the filter fields to turn on the locator lights for the displayed drives.

This action helps you physically locate the drives in the storage array.

Turn on drive locator light

From the Hardware page, you can turn on the locator light to find the physical location of a drive in the storage array.

About this task

You can locate single drives or multiple drives shown on the Hardware page.

Steps

1. Select **Hardware**.

2. To locate one or more drives, do one of the following:

- **Single drive** — From the shelf graphic, find the drive you want to physically locate in the array. (If the graphic shows the controllers, click **Show front of shelf**.) Click the drive to display its context menu, and then select **Turn on locator light**.

  The drive's locator light turns on. When you have physically located the drive, return to the dialog and select **Turn off**.

- **Multiple drives** — In the filter fields, select a physical drive type from the left drop-down list and a logical drive type from the right drop-down list. The number of drives matching your criteria is shown on the far right of the fields. Next, you can either click **Turn on locator lights** or select **Locate all filtered drives** from the context menu. When you have physically located the drives, return to the dialog and select **Turn off**.

View drive status and settings

You can view status and settings for the drives, such as the media type, interface type, and capacity.

Steps

1. Select **Hardware**.

2. If the graphic shows the controllers, click **Show front of shelf**.

   The graphic changes to show the drives instead of the controllers.

3. Select the drive for which you want to view status and settings.

   The drive's context menu opens.

4. Select **View settings**.

   The Drive Settings dialog box opens.

5. To see all settings, click **Show more settings**, in the upper right of the dialog box.

Field Details

Settings	Description
Status	Displays Optimal, Offline, Non-critical fault, and Failed. Optimal status indicates the desired working condition.
Mode	Displays Assigned, Unassigned, Hot Spare Standby, or Hot Spare in Use.
Location	Shows the shelf and bay number where the drive is located.
Assigned to/Can protect for/Protecting	If the drive is assigned to a pool, volume group, or SSD cache, this field displays "Assigned to." The value can be a pool name, volume group name, or SSD cache name. If the drive is assigned to a hot spare and its mode is Standby, this field displays "Can protect for." If the hot spare can protect one or more volume groups, the volume group names appear. If it cannot protect a volume group, it displays 0 volume groups.

If the drive is assigned to a hot spare and its mode is In Use, this field displays "Protecting." The value is the name of the affected volume group.

If the drive is unassigned, this field does not appear. |
Media type	Displays the type of recording media the drive uses, which can be either hard disk drive (HDD) or solid state disk (SSD).
Percent endurance used (only shown if SSD drives are present)	The amount of data written to the drive to date, divided by the total theoretical write limit.
Interface type	Displays the type of interface the drive uses, such as SAS.
Drive path redundancy	Shows whether connections between the drive and controller are redundant (Yes) or not (No).
Capacity (GiB)	Shows the usable capacity (total configured capacity) of the drive.
Speed (RPM)	Shows the speed in RPM (does not appear for SSDs).
Current data rate	Shows the data transfer rate between the drive and the storage array.
Logical sector size (bytes)	Shows the logical sector size that the drive uses.
Physical sector size (bytes)	Shows the physical sector size that the drive uses. Typically, the physical sector size is 4096 bytes for hard disk drives.
Drive firmware version	Shows the revision level of the drive firmware.

Settings	Description
World-wide identifier	Shows the unique hexadecimal identifier for the drive.
Product ID	Shows the product identifier, which is assigned by the manufacturer.
Serial number	Shows the serial number of the drive.
Manufacturer	Shows the vendor of the drive.
Date of manufacture	Shows the date the drive was built. (i) Not available for NVMe drives.
Secure-capable	Shows whether the drive is secure-capable (Yes) or not (No). Secure-capable drives can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives, which encrypt data during writes and decrypt data during reads. These drives are considered secure-*capable* because they can be used for additional security using the Drive Security feature. If the Drive Security feature is enabled for volume groups and pools used with these drives, the drives become secure-*enabled*.
Secure-enabled	Shows whether the drive is secure-enabled (Yes) or not (No). Secure-enabled drives are used with the Drive Security feature. When you enable the Drive Security feature and then apply Drive Security to a pool or volume group on secure-*capable* drives, the drives become secure-*enabled*. Read and write access is available only through a controller that is configured with the correct security key. This added security prevents unauthorized access to the data on a drive that is physically removed from the storage array.
Data Assurance (DA) capable	Shows whether the Data Assurance (DA) feature is enabled (Yes) or not (No). Data Assurance (DA) is a feature that checks for and corrects errors that might occur as data is transferred through the controllers down to the drives. Data Assurance can be enabled at the pool or volume group level, with hosts using a DA-capable I/O interface such as Fibre Channel.
Read/write accessible	Shows whether the drive is read/write accessible (Yes) or not (No).
Drive security key identifier	Shows the security key for secure-enabled drives. Drive Security is a storage array feature that provides an extra layer of security with either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives. When these drives are used with the Drive Security feature, they require a security key for access to their data. When the drives are physically removed from the array, they cannot operate until they are installed in another array, at which point, they will be in a Security Locked state until the correct security key is provided.

6. Click **Close**.

Replace drive logically

If a drive fails or you want to replace it for any other reason, and you have an unassigned drive in your storage array, you can logically replace the failed drive with the unassigned drive. If you do not have an unassigned drive, you can physically replace the drive instead.

About this task

When you logically replace a drive with an unassigned drive, the unassigned drive becomes assigned and is then a permanent member of the associated pool or volume group. You use the logical replace option to replace the following types of drives:

- Failed drives
- Missing drives
- SSD drives that the Recovery Guru has notified you that are nearing their end of life
- Hard drives that the Recovery Guru has notified you that have an impending drive failure
- Assigned drives (available only for drives in a volume group, not in a pool)

The replacement drive must have the following characteristics:

- In the Optimal state
- In the Unassigned state
- The same attributes as the drive being replaced (media type, interface type, and so on)
- The same FDE capability (recommended, but not required)
- The same DA capability (recommended, but not required)

Steps

1. Select **Hardware**.

2. If the graphic shows the controllers, click **Show front of shelf**.

   The graphic changes to show the drives instead of the controllers.

3. Click the drive that you want to logically replace.

   The drive's context menu appears.

4. Click **Logically replace**.

5. **Optional:** Select the **Fail drive after it is replaced** check box to fail the original drive after it is replaced.

   This check box is enabled only if the original assigned drive is not failed or missing.

6. From the **Select a replacement drive** table, select the replacement drive that you want to use.

   The table lists only those drives that are compatible with the drive that you are replacing. If possible, select a drive that will maintain shelf loss protection and drawer loss protection.

7. Click **Replace**.

If the original drive is failed or missing, data is reconstructed on the replacement drive using the parity information. This reconstruction begins automatically. The drive's fault indicator lights go off, and the activity indicator lights of the drives in the pool or volume group start flashing.

If the original drive is not failed or missing, its data is copied to the replacement drive. This copy operation begins automatically. After the copy operation completes, the system transitions the original drive to the Unassigned state, or if the check box was selected, to the Failed state.

Reconstruct drive manually

Drive reconstruction normally starts automatically after you replace a drive. If drive reconstruction does not start automatically, you can start reconstruction manually.

About this task

ⓘ   Perform this operation only when instructed to do so by technical support or the Recovery Guru.

Steps

1. Select **Hardware**.
2. If the graphic shows the controllers, click **Show front of shelf**.

   The graphic changes to show the drives instead of the controllers.

3. Click the drive that you want to manually reconstruct.

   The drive's context menu appears.

4. Select **Reconstruct**, and confirm that you want to perform the operation.

Initialize (format) drive

If you move assigned drives from one storage array to another, you must initialize (format) the drives before they can be used in the new storage array.

About this task

Initializing removes the previous configuration information from a drive and returns it to the Unassigned state. The drive is then available for adding to a new pool or volume group in the new storage array.

Use the initialize drive operation when you are moving a single drive. You do not need to initialize drives if you are moving an entire volume group from one storage array to another.

⚠   **Possible loss of data** — When you initialize a drive, all data on the drive is lost. Perform this operation only when instructed to do so by technical support.

Steps

1. Select **Hardware**.
2. If the graphic shows the controllers, click **Show front of shelf**.

   The graphic changes to show the drives instead of the controllers.

3. Click the drive that you want to initialize.

   The drive's context menu appears.

4. Select **Initialize**, and confirm that you want to perform the operation.

Fail drive

If instructed to do so, you can manually fail a drive.

About this task

System Manager monitors the drives in the storage array. When it detects that a drive is generating a lot of errors, the Recovery Guru notifies you of an impending drive failure. If this happens and you have a replacement drive available, you might want to fail the drive to take preemptive action. If you do not have a replacement drive available, you can wait for the drive to fail on its own.

> ⚠️ **Possible loss of data access** — This operation could result in data loss or the loss of data redundancy. Perform this operation only when instructed to do so by technical support or the Recovery Guru.

Steps

1. Select **Hardware**.

2. If the graphic shows the controllers, click **Show front of shelf**.

   The graphic changes to show the drives instead of the controllers.

3. Click the drive that you want to fail.

   The drive's context menu appears.

4. Select **Fail**.

5. Keep the **Copy contents of drive before failing** check box selected.

   The copy option appears only for assigned drives and for non-RAID 0 volume groups.

   Before you fail the drive, make sure that you copy the drive's contents. Depending on your configuration, you could potentially lose all data or data redundancy on the associated pool or volume group if you do not copy the drive's contents first.

   The copy option allows faster drive recovery than reconstruction and reduces the possibility of a volume failure if another drive were to fail during the copy operation.

6. Confirm that you want to fail the drive.

   After the drive has failed, wait at least 30 seconds before you remove it.

Assign hot spares

You can assign a hot spare as a standby drive for additional data protection in RAID 1, RAID 5, or RAID 6 volume groups. If a drive fails in one of these volume groups, the controller reconstructs data from the failed drive to the hot spare.

Before you begin

- RAID 1, RAID 5, or RAID 6 volume groups must be created. (Hot spares cannot be used for pools. Instead, a pool uses spare capacity within each drive for its data protection.)

- A drive that meets the following criteria must be available:
  - Unassigned, with Optimal status.
  - Same media type as the drives in the volume group (for example, SSDs).
  - Same interface type as the drives in the volume group (for example, SAS).
  - Capacity equal to or larger than the used capacity of the drives in the volume group.

About this task

This task describes how to manually assign a hot spare from the Hardware page. The recommended coverage is two hot spares per drive set.

> ⓘ Hot spares can also be assigned from the Initial Setup wizard. You can determine if hot spares are already assigned by looking for drive bays shown in pink on the Hardware page.

Steps

1. Select **Hardware**.

2. If the graphic shows the controllers, click **Show front of shelf**.

   The graphic changes to show the drives instead of the controllers.

3. Select an unassigned drive (shown in gray) that you want to use as a hot spare.

   The drive's context menu opens.

4. Select **Assign hot spare**.

   If the drive is secure-enabled, the Secure Erase Drive? dialog box opens. To use a secure-enabled drive as a hot spare, you must first perform a Secure Erase operation to remove all its data and reset its security attributes.

   > ⚠ **Possible loss of data** — Make sure that you have selected the correct drive. After completing the Secure Erase operation, you cannot recover any of the data.

   If the drive is **not** secure-enabled, the Confirm Assign Hot Spare Drive dialog box opens.

5. Review the text in the dialog box, and then confirm the operation.

   The drive is displayed in pink on the Hardware page, which indicates it is now a hot spare.

Results

If a drive within a RAID 1, RAID 5, or RAID 6 volume group fails, the controller automatically uses redundancy data to reconstruct the data from the failed drive to the hot spare.

Unassign hot spares

You can change a hot spare back to an unassigned drive.

Before you begin

The hot spare must be in Optimal, Standby status.

About this task

You cannot unassign a hot spare that is currently taking over for a failed drive. If the hot spare is not in Optimal status, follow the Recovery Guru procedures to correct any problems before trying to unassign the drive.

Steps

1. Select **Hardware**.

2. If the graphic shows the controllers, click **Show front of shelf**.

   The graphic changes to show the drives instead of the controllers.

3. Select the hot spare drive (displayed in pink) that you want to unassign.

   If there are diagonal lines through the pink drive bay, the hot spare is currently in use and cannot be unassigned.

   The drive's context menu opens.

4. From the drive's drop-down list, select **Unassign hot spare**.

   The dialog box shows any volume groups affected by removing this hot spare and if any other hot spares are protecting them.

5. Confirm the unassign operation.

Results

The drive is returned to Unassigned (shown in gray).

Erase secure-enabled drive

You can erase a secure-enabled drive so it can be reused in another volume group, pool, SSD cache, or in another storage array. This procedure resets the drive's security attributes and ensures that the data cannot be read again.

Before you begin

The secure-enabled drive must be in an Unassigned state.

About this task

Use the Secure Erase option only if you want to remove all data on a secure-enabled drive and reset the drive's security attributes.

> ⚠ **Possible loss of data** — The Secure Erase operation cannot be undone. Make sure you select the correct drive during the procedure.

Steps

1. Select **Hardware**.

2. If the graphic shows the controllers, click **Show front of shelf**.

   The graphic changes to show the drives instead of the controllers.

3. Use the filter fields to view all the secure-enabled, unassigned drives in the shelf. From the **Show drives that are…** drop-down lists, select **Secure-enabled** and **Unassigned**.

> (i) If all drives share the same physical attributes, the **Show drives that are…** filter field does not appear. If all drives share the same logical attributes, the **Anywhere in the storage array** filter field does not appear.
>
> The shelf view shows only the secure-enabled, unassigned drives; all others are grayed out.

4. Select the secure-enabled drive you want to erase.

> (!) **Possible loss of data** — Make sure that you have selected the correct drive. After completing the Secure Erase operation, you cannot recover any of the data.

The drive's context menu opens.

5. Select **Secure Erase**.

The Secure Erase option only appears if you select an unassigned, secure-enabled drive.

> (i) For NVMe SED drives, you must provide the PSID. You can find the PSID on the drive label. This is necessary if you do not have the backup lock key.

6. In the Secure Erase Drive dialog box, read the important information about data loss.

7. Confirm the operation, and then click **Erase**.

Results

The drive is now available for use in another volume group or disk pool, or in another storage array.

Unlock or reset locked NVMe drives

If you insert one or more locked NVMe drives into a storage array, you can unlock the drive data by adding the security key file associated with the drives. If you do not have a security key, you can perform a reset on each locked NVMe drive by entering its Physical Security ID (PSID) to reset its security attributes and erase the drive data.

Before you begin

- For the Unlock option, make sure the security key file (with an extension of `.slk`) is available on the management client (the system with a browser used for accessing System Manager). You must also know the pass phrase associated with the key.

- For the Reset option, you must find the PSID on each drive you want to reset. To locate the PSID, physically remove the drive and locate the PSID string (32 characters maximum) on the drive's label, and then reinstall the drive.

About this task

This task describes how to unlock data in NVMe drives by importing a security key file into the storage array. For situations where the security key is not available, this task also describes how to perform a reset on a locked drive.

> ⓘ If the drive was locked using an external key management server, select **Settings › System › Security key management** in System Manager to configure external key management and unlock the drive.

Steps

1. Select **Hardware**.

2. If the graphic shows the controllers, click **Show front of shelf**.

   The graphic changes to show the drives instead of the controllers.

3. Select the NVMe drive you want to unlock or reset.

   The drive's context menu opens.

4. Select **Unlock** to apply the security key file or **Reset** if you do not have a security key file.

   These options only appear if you select a locked NVMe drive.

   > ⚠ During a Reset operation, all data is erased. Only perform a Reset if you do not have a security key. Resetting a locked drive permanently removes all data on the drive and resets its security attributes to "secure-capable," but not enabled. **This operation is not reversible.**

5. Do one of the following:

   a. **Unlock**: In the Unlock Secure Drive dialog box, click **Browse**, and then select the security key file that corresponds to the drive you want to unlock. Next, enter the pass phrase, and then click **Unlock**.

   b. **Reset**: In the Reset Locked Drive dialog box, enter the PSID string in the field, and then type `RESET` to confirm. Click **Reset**.

   For an Unlock operation, you only need to perform this operation once to unlock all the NVMe drives. For a Reset operation, you must individually select each drive you want to reset.

Results

The NVMe drive is now available for use in another volume group or disk pool, or in another storage array.

FAQs

What is preservation capacity?

Preservation capacity is the amount of capacity (number of drives) that is reserved in a pool to support potential drive failures.

When a pool is created, System Manager automatically reserves a default amount of preservation capacity depending on the number of drives in the pool.

Pools use preservation capacity during reconstruction, whereas volume groups use hot spare drives for the same purpose. The preservation capacity method is an improvement over hot spare drives because it allows reconstruction to happen faster. Preservation capacity is spread over a number of drives in the pool instead of on one drive in the case of a hot spare drive, so you are not limited by the speed or availability of one drive.

Why would I logically replace a drive?

If a drive fails or you want to replace it for any other reason, and you have an unassigned drive in your storage array, you can logically replace the failed drive with the unassigned drive. If you do not have an unassigned drive, you can physically replace the drive instead.

The data from the original drive is copied or reconstructed onto the replacement drive.

Where can I view the status of a drive undergoing reconstruction?

You can view drive reconstruction status from the Operations in Progress dashboard.

From the **Home** page, click the **View Operations in Progress** link in the upper right.

Depending on the drive, the full reconstruction might take a considerable amount of time. If a volume ownership has changed, a full reconstruction might take place instead of the rapid reconstruction.